# PCI Express External Cabling Specification

Revision 3.0, Version 0.7

November 9, 2016

PCI-SIG® disclaims all warranties and liability for the use of this document and the information contained herein and assumes no responsibility for any errors that appear in this document, nor does PCI-SIG make a commitment to update the information contained herein.

Contact the PCI-SIG office to obtain the latest revision of this specification.

Forward ant questions regarding the PCI Code and ID Assignment Specification or membership in PCI-SIG to:

**Membership Services**
www.pcisig.com
E-mail: administration@pcisig.com
Phone: 503-619-0569
Fax: 503-644-6708

**Technical Support**
techsupp@pcisig.com

**DISCLAIMER**
This PCI Code and ID Assignment Specification is provided *as is* with no warranties whatsoever, including any warranty of merchantability, non-infringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification, or sample. PCI-SIG disclaims all liability for infringement of proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

**Trademarks**
PCI, PCI Express, PCIe, ExpressModule, and PCI-SIG are trademarks or registered trademarks of PCI SIG. Other company and product names are trademarks of the respective companies with which they are associated.

# Revision History

| Revision | Version | History | Date |
|---|---|---|---|
| 1.0 | | Initial public release | January 7, 2007 |
| 2.0 | | Release includes 5.0GT/s | July 22, 2012 |
| 3.0 | 0.1 | Initial working draft | June 19, 2013 |
| | 0.3 | Updated revision, updated header | October 10, 2013 |
| | 0.5 | Following additions/updates were done:<br>• Added MGTPWR<br>• Added port configuration options<br>• Added sideband function<br>• Added SAS interop (preliminary)<br>• Added notes to the memory map, the new style guide is currently not being referenced,<br>• Updated 2-wire interface<br>• Updated diagrams<br>• Changed SMBus to CMI<br>• Swapped pin assignments for lanes 0 & 1<br>• Removed the cable clock from the Upstream system completely<br>• Moved CREFCLK information to Appendix for Legacy Adapter Cable<br>• Updated images<br>• Updated terms and punctuation for consistency | March 3, 2015 |
| | 0.6 | • Included feedback from 0.5 feedback<br>• Updated some pinout tables with correct names<br>• Removed notes column from cable memory map tables<br>• Added power sequencing requirements<br>• Updated cable performance requirements, included recommended settings for RxEQ based on cable loss | May 25, 2015 |
| | 0.7 | Updated the following:<br>• 6.2.2.3.5, Bytes 108-109, Propagation Delay<br>• Added Figure 6-3, Bytes 108-109, Propagation Delay<br>• Added Table 6-6, Lower Bound, Upper Bound, Base, and Span Values<br>• Updated equations in Chapter 6, Cable Specification<br>• Updated language per PCI SIG Style Guide | November 9, 2016 |
| | | | |

# Table of Contents

# List of Figures

# List of Tables

193

1

194 # 1. Introduction

195 This is a companion specification to the *PCI Express Base Specification*. The primary focus of this
196 specification is the implementation of cabled PCI Express®. No assumptions are made regarding the
197 implementation of PCI Express-compliant Subsystems on either side of the cabled Link (PCI
198 ExpressCard Electromechanical (CEM), ExpressCard™, ExpressModule™, PXI Express™, or any
199 other form factor). Such form factors are covered in separate specifications.

200 ## 1.1. Terms and Acronyms

| | |
|---|---|
| x1, x2, x4, x8, x12, x16 | x1 refers to one PCI Express Lane of Basic bandwidth; x4 to a collection of four PCI Express Lanes; etc. |
| Active Cable Assembly | A cable assembly that uses active components to buffer or condition the PCI Express signals. |
| Auxiliary signals | Signals not required by the PCI Express architecture, but necessary for certain desired functions or system implementation; for example, the CBLPRSNT# signal. |
| Cable Management Interface, CMI | The 2-wire management interface for communication with the cable and between Subsystems. |
| Cable Aggregation | The ability of a port to combine multiple cable assemblies inserted to the external connector to create a port larger than any of the individual connectors. |
| Cable Link Partner, CLP | The Fixed-Side Subsystem that a device and port are connected to via a cable |
| Device | A component on either end of a PCI Express Link. |
| Downstream | 1. The relative position of an inter-connect/system element (Port/component) that is farther from the Root Complex. Within the context of this specification also referred to as *Downstream Subsystem*.<br>2. A direction of information flow where the information is flowing *away* from the Root Complex. |
| Downstream Facing Port, DSF | A cabled port in the Upstream Subsystem that links over the cable to a Downstream Subsystem. |

| | |
|---|---|
| Endpoint | A device with a Type 00h Configuration Space header. |
| Fixed-Side | The Subsystem a cable plugs into. This references either the Upstream or Downstream cable port; the term itself does not distinguish which end of the topology. |
| Free-Side | This is the cable side of a port. It is the removable side, or free side, of a cable connection. |
| Form Factor | In the context of this specification, form factor refers to other specifications that could be used as the foundation for implementing an external cabled PCI Express Port; e.g., PCI ExpressCard Electromechanical, ExpressCard, ExpressModule, Compact PCI Express. |
| G3 | Global Power State in which no external power is provided to the Box. Refer to the ACPI specification, Revision 5.0. |
| HCSL | High-speed Current Steering Logic implemented in various different clock generator components. |
| Host System | The compute entity which contains the PCI Express Root Complex and is the source of the reference clock signal. In the context of this specification, also referred to as *Upstream Subsystem*. |
| Hot-Plug | Insertion and/or removal of a cable and Downstream Subsystem into/from an active Upstream Subsystem. |
| Lane | One PCI Express Lane contains a differential pair for transmission and another differential pair for reception. A by-N, or xN, Link is composed of N Lanes. |
| Link | A collection of one or more PCI Express Lanes providing the communication path between an Upstream and Downstream Port. |
| Legacy | Reference to the cable and parameters specified in the first two revisions of this specification that support 2.5 GT/s and 5.0 GT/s. |
| Local | A Fixed-Side port and inserted cable connector. |
| Multi-master Controller | A device that implements the Multi-master functions, with arbitration, specified in the *I2C Specification*. |
| Passive Cable Assembly | A cable assembly that does not use any active components for the PCI Express signals. The cable assembly must still contain a memory device for system configuration information. |
| Port | 1. A group of Transmitters and Receivers located on the same device that define a Link when active.<br>2. Physical connector for a Fixed-Subsystem related to a link, typically referring to the Fixed-side |
| Port Bifurcation | The ability of an Upstream Subsystem to configure an external connector as multiple smaller ports.. |
| Sideband Messages | A method for signaling events and conditions using the 2-wire interface between two linked components. |
| Subsystem | In the context of this specification, Subsystem is a generic term identifying either an Upstream or Downstream Fixed-Side containing the PCI-Express device, CMI controller, and other circuits providing a cabled PCI Express Port. |

| | |
|---|---|
| Upstream | 1. The relative position of interconnect/system element (Port/component) that is nearest to the Root Complex. Within the context of this specification also referred to as *Upstream Subsystem*.<br>2. A direction of information flow where the information is flowing toward the Root Complex. |
| Upstream Facing Port, USF | A cabled port in the Downstream Subsystem that links over the cable to an Upstream Subsystem. |
| Vendor Specific Message | A message sent across the cable over the Cable Management Interface that is defined by the vendor. This is not related to in-band vendor-defined Messages allowed in the protocol. |
| Wake | A mechanism used by a component to request the reapplication of main power when in the L2 Link state. Two such mechanisms are defined in the *PCI Express Base Specification*: Beacon and WAKE#. |

## 1.2. Reference Documents

Following is a list of support documentation that should be referenced:

- *PCI Express Base Specification*, Revision 3.1 (with SRIS ECN)
- *PCI ExpressCard Electromechanical Specification*, Revision 3.0
- *PCI Hot-Plug Specification*, Rev. 1.1
- *PCI Standard Hot-Plug Controller and Subsystem Specification*, Rev. 1.0
- EIA 364 Series, Electrical Connector Test Procedures Including Environmental Classifications with Test Procedures
- EIA 364-1000, Environmental Test Methodology for Assessing the Performance of Connectors and Sockets used in Business Office Applications
- *SFF-8644, Mini Multilane 12 Gbs 8/4x Shielded Connector*, Rev. 2.9
- *SFF-8636, Common Management Interface*, Rev. 2.6
- *I2C-Bus Specification* Revision 2.1
- *Advanced Configuration and Power Interface Specification* 5.0
- *SFF-8449, Shielded Cables Management Interface for SAS, Rev 2.0.*

# 1.3. Documentation Conventions

Following are a list of conventions that have been used throughout this specification:

- **Capitalization**
  Some terms are capitalized to distinguish their definition in the context of this document from their common English meaning Words not capitalized have their common English meaning. When terms such as *memory write* or *memory read* appear completely in lower case, they include all transactions of that type.
  Register names, and the names of fields and bits in registers and headers, are presented with the first letter capitalized and the remainder in lower case.
- **Numbers and Number Bases**
  Hexadecimal numbers are written with a lower case "h" suffix (like FFFh and 80h). Hexadecimal numbers larger than four digits are represented with a space dividing each group of four digits, as in 1E FFFF FFFFh. Binary numbers are written with a lower case "b" suffix (1001b and 10b). Binary numbers larger than four digits are written with a space dividing each group of four digits, as in 1000 0101 0010b.
  All other numbers are decimal.
- **Implementation Notes**
  Implementation Notes should not be considered to be part of this specification. They are included for clarification and illustration only.

# 1.4. Specification Contents

This specification contains the following information:

- Subsystem requirements
- Sideband signaling and usage models
- Cable hot insertion and removal
- Subsystem electrical budgets
- Cable and connectors electrical budget
- Cable and connectors specifications
- Power provisioning

This specification is broken up into three primary sections. The first discusses the new Third Generation of the external cabling interface specifications and covers the new connectors, cables, signaling, Cable Management Interface, and sidebands. The second section discusses the previous revision of the specification that covers the interface defined for the First and Second Generation of this specification. The third section contains appendices that discusses various other topics related to the implementations of cabled interfaces.

## 1.5. Objectives

The objectives of this specification are as follows:

- Define PCI Express external cables and associated connectors
- Support PCI Express data rates of up to 8.0 GT/s, while maintaining compatibility with the 5.0 GT/s and 2.5 GT/s signaling rate specifications
- Support standard PCI Express components, as defined by the PCI Express Base Specification, Revision 3.0
- Forward looking for future scalability
- Maximize cable interoperability for user flexibility
- Enable Hot-plug as a native function
- Allow revolutionary partitioning of the PC architecture
- Upgradeability
- Backwards compatibility with Subsystems compatible with previous versions of this specification

## 1.6. Overview

PCI Express is the third generation of a multi-purpose I/O interface that are used across the computing industry (from mobile through high-end servers and communication equipment.) The broad usage and versatility of this technology allows for system extensions to external input/output Subsystems that meet the needs for specific target applications and/or environments.

- Cabled PCI Express targets a large number of applications, including but not limited to:
- Split-systems, or disaggregate PCs, with a desktop "console" that contains removable media drives (e.g., CD/DVD), memory modules, I/O ports (e.g., USB, IEEE-1394), and audio jacks
- I/O expansion to extend the I/O card capabilities of the main system for support of different form factors, including legacy Subsystems, test and measurement, and instrumentation equipment
- Server expansion I/O to support conventional PCI Express add-in cards (with or without Hot-Plug support) and/or ExpressModules
- Location of the graphic Subsystem (i.e., controller and memory) external to the main systems unit

System-level support for cabled PCI Express is possible through implementation on expansion card (on a PCI Express CEM, ExpressModules, PXI Express, ExpressCard, or direct from a system board).

This specification defines connectors and implementation for supporting up to x16 cabled PCI Express Links.

Utilizing off-the-shelf PCI Express components is the focus of this specification, while providing enough flexibility for design and implementation of dedicated components for driving the cabled interconnect. As such, much of this specification builds upon other PCI Express form factor specifications. However, not all components inherently function correctly in this application.

New to this version of the specification are electrical and mechanical requirements. These requirements may introduce some backward compatibility concerns.

**Note:** This revision of the specification calls for a new connector family, one commonly known as *SFF-8644*, that is not mechanically compatible with the connectors of the previous versions of this specification. Therefore, cable ports and assemblies designed to this specification are not directly compatible with ports and assemblies designed specifically to meet the previous revisions of this specification. The new connector scheme provides for aggregation of connectors and cable assemblies to support Links widths of x16, as well as the option for Port Bifurcation, to support multiple links to a single Upstream physical port/connector. See Figure 1-1.

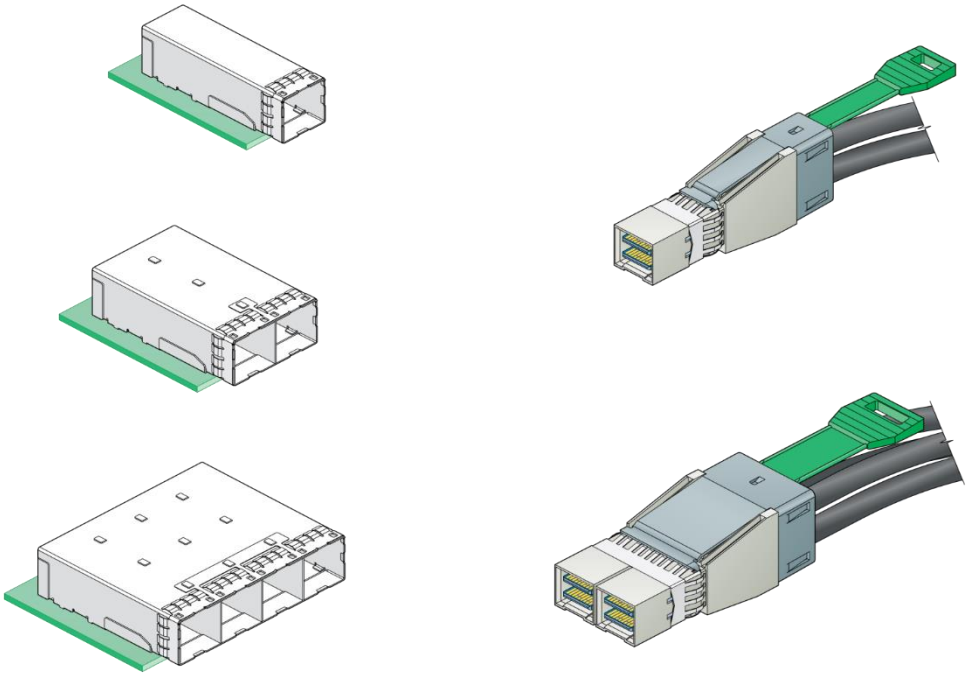

**Figure 1-1. Connector for Fixed-Side and Free-Side of Cabled Interface**

New features enabled by this revision of the specification include:

- Bifurcation of cable ports, which enables Upstream facing ports with more than one x4 connector to be split into smaller, separate links. (Optional)
- Cable Aggregation, which allows multiple smaller cable widths to create a larger port. (Optional)
- Sideband Messages, a new mechanism to send sideband signals over a serial protocol.
- Vendor Specific Communication, a mechanism for sending Vendor Specific Messages outside of the PCI-Express signal path. (Optional)
- Cable Monitoring, which allows for monitoring of active cable component status. (Optional)
- Fiber Optic Support (Optional)
- Compatibility with SFF-8449 cables. These cables are be used to create x4 links and do not have access to several of the above features.

Previous versions of this specification utilized/employed a fixed-size component for each supported link width (e.g., a x1 link had specific x1 cable ports and a x1 cable assembly, a x4 link had a specific x4 cable port and x4 specific cable assembly, up through x16 link widths). Unlike some other form-factors, components of various link widths were not interoperable. Introduced in this revision of this specification, the SFF-8644 connectors are defined for increased performance characteristics and new features. The connector family scales a link from x4 connection up to x16. A unique x1 connector is not provided and instead the new x4 connector is used for any size link up to, and including, x4 link. The x8 link aggregates two x4 links, and x16 aggregates four x4 links. This allows for a more universal and scalable cable port.

The introduction of 8.0 GT/s in the *PCI Express Base Specification*, requires some foreknowledge of the physical link interconnect electrical performance for the purpose of configuring equalizer starting points. Since the performance of a cable is unknown to either Fixed-Side Subsystem, the use of the CMI and cables with memory has been employed for storing and retrieving cable assembly performance information (Figure 1-2 shows the cable interface).

The sideband messaging scheme changes such that individual wires are no longer used for each sideband signal. This revision of the specification also enables an optical friendly architecture for those needing either distance or electrical isolation between Subsystems.

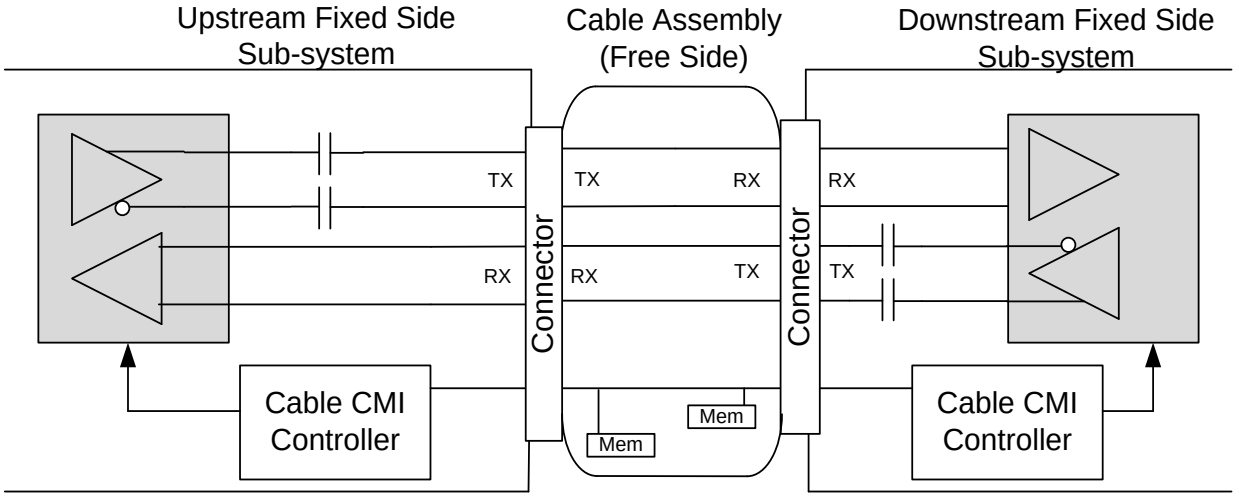


**Figure 1-2. Cable Interface**

2

# 2. Auxiliary Signals

Auxiliary signals are provided on the connector to assist with system-level functionality or implementation. The high-speed differential signaling levels are compatible with advanced silicon processes while all low-speed Auxiliary signals are defined to be compatible with +3.3 V signaling unless otherwise noted.

Besides the signals that are required to transmit/receive data on the PCI Express interface, there are also signals that may be necessary to implement the PCI Express interface in a distributed system environment, or to provide certain desired functions, including power for data conditioning within the connector backshell. These signals are referred to as the auxiliary or sideband signals.

The PCI Express cable connector and cabling support the following Auxiliary signals:

- **PWR**: Power provisioning to the connector backshell is provided to allow for signal conditioning components within the cable assembly. A wire must not be provided within the cable.
- **MGTPWR**: Power supplied to the connector backshell for cable management components that are needed while the link is not active. This needs to be active if the Subsystem has power. A wire must not be provided within the cable.
- **CBLPRSNT#**: Cable present detect, an active-low signal pulled-down by the Free-Side when it is inserted into the Fixed-Side Connector.
- **CADDR**: This signal is used to configure the Upstream cable management device address.
- **CINT#**: This signal is asserted by the cable assembly to indicate a need for service via the Cable Management Interface controller.
- **CMISDA**: Management interface data line. Used for both initial link setup and sideband messages when used with proper cables.
- **CMISCL**: Management interface clock line. Used for both initial link setup and sideband messages when used with proper cables.

## 2.1. PWR

Power for the optional active cable circuitry within the cable plug is required. This pin may not be powered in all Subsystem power states such as when the Subsystem is in a power saving state.

Due to the low level signaling of the PCI Express interface, it is strongly recommended that sufficient decoupling of all power supplies be provided. This is recommended to ensure that power supply noise does not interfere with the recovery of data from a remote Subsystem. Some basic guidelines to help ensure a quiet power supply are provided section 2.9.

**Note:** The details given in this section are guidelines only. It is the responsibility of the designer to properly test the design to ensure that system board circuitry does not create excessive noise on power supply or ground signals at the connector interface.

To protect system operation, the PCI Express external cable assemblies (that operate from board-supplied power) must follow the requirements listed in Table 2-3 during Hot-Plugging and normal operation. **Error! Reference source not found.** shows the power requirements. The requirements f or the power supply voltages are defined at the fixed-side connector.

## 2.2. MGTPWR

MGTPWR to the cable assembly is required to enable cable management circuitry. It is recommended that MGTPWR remain within operational limits when in power states other than G3 to enable remote power control and decrease cabled link enable times.

**Note:** MGTPWR is removed and re-applied to reset the cable logic.

## 2.3. CBLPRSNT#

A cable assembly presence detection mechanism is specified, through CBLPRSNT#, to indicate the attachment of a cable assembly to a Fixed-Side port. The Free-Side must assert this signal low to indicate cable attachment. It is recommended that this signal not be directly grounded, but rather asserted with a pulldown or active driver. The assertion of CBLPRSNT# provides no information about the far end of the cable, including whether or not it is inserted into the Fixed-Side of a CLP.

Upon assertion of CBLPRSNT#, the Fixed-Side management controller must read the cable memory's first 256 bytes to obtain cable information and begin monitoring of CINT# Refer to 4.2.1 for specific usage in power sequencing. If the cable supports sideband messages, as indicated by the Support for Sideband Messages bit in the Cabled PCIe Capabilities 1 register, the management controller is permitted to attempt to detect the presence of the CLP by reading the opposite Fixed-Side via CMI. This acts as a Hot-Plug detection scheme.

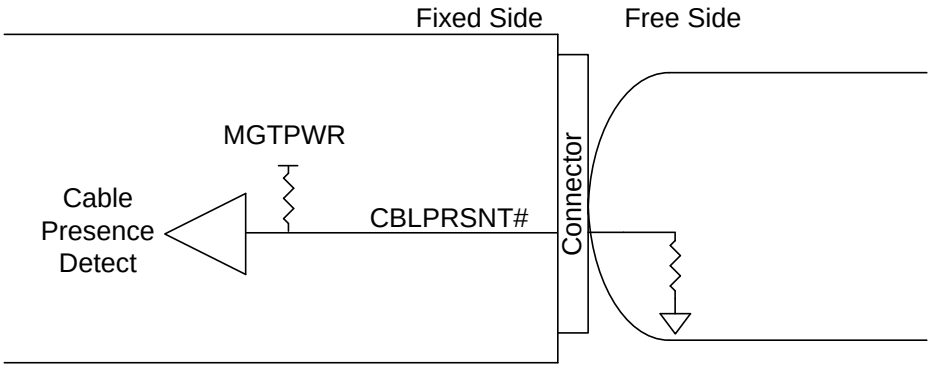


**Figure 2-1. CBLPRSNT# Signaling Diagram**

> **Observation:** CBLPRSNT# being asserted by the cable determines if the cable is present. This is used as a link disable by the cable assembly by deasserting CBLPRSNT# in an active fashion. This may be useful for some cable assemblies, but specific implementations are beyond the scope of this specification.

# 2.4. CADDR

The CADDR signal enables the use of a programmable address pin for the Free-Side CMI. Cable assemblies using SFF-8644 connectors at both ends use this pin as the address reference for the Least Significant Bit, or LSB, of the Free-Side memory address. The Free-Side must tie the LSB of the 7-bit address low with a resistor that allows the signal to remain below $V_{OL\_SB}$ when the Fixed-Side is open and above $V_{OH\_SB}$ when the Fixed-Side has a pull-up of $R_{PU\_AUX}$ at both ends of the cable assembly. The address may be changed by the Fixed-Side by overdriving the weak pull-down of the Free-Side. For cable assemblies supporting side-band messages, only address A0h must contain the cable description information, the duplicate information at A2h is optional. If there is a single memory device, CADDR may be ignored within the cable assembly such that the address is fixed to A0h. However, cable assembly status information, such as the interrupts and monitors described in paragraph 6.2.1, are only be available to the USF Fixed-Side. Upstream Fixed-side management controllers must be read from address A0h for configuration data if *acknowledge* is not received when attempting to read from address A2h.

**Observation:** Upstream fixed management engine may listen for configuration information at address A0h and the Downstream fixed management engine may listen to configuration information at address A2h. This may decrease configuration time.

The Upstream Fixed-Side must drive CADDR high, via resistor of value $R_{PU\_AUX}$, or active driver.

Support of a cross-link configuration requires an active driver on the Upstream Fixed-Side. Once the PCI-Express protocol has determined which portion of the link is the Upstream Subsystem, the Fixed-Side Subsystem that has been determined to be the Downstream Subsystem must allow the CADDR to float to configure the Downstream end of the cable for management purposes. In the cross-link system, beginning with initialization, there are times when both cable memory locations are at the same address. Upstream Subsystems that do not support crosslink are permitted to implement a passive solution.

The Downstream Fixed-Side must not connect to CADDR. The Downstream end of the cable uses the default address of A0h for the cable memory with CADDR not connected by the Fixed-side.

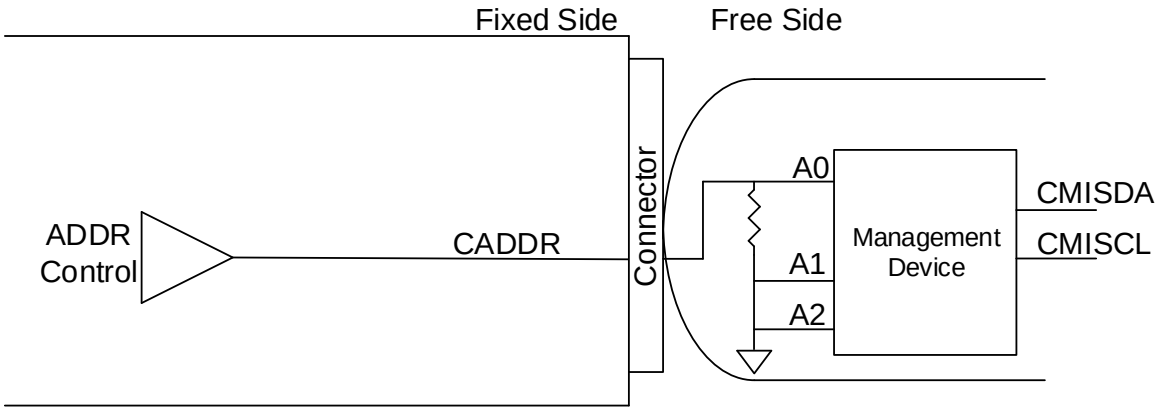


**Figure 2-2. CADDR Signal Diagram**

# 2.5. CINT#

CINT# enables the use of a Free-Side driven interrupt to indicate a need for service via the Cable Management Interface. CINT# is localized to the near end Free-Side of the cable only and does not provide any indication across the cable to the opposite Fixed-Side. The Upstream Fixed-Side receives only the CINT# from the end of the cable inserted to the Upstream Free-Side and then reads address A2h, and the Downstream Fixed-Side receives only the CINT# from the Downstream Free-Side and reads address A0h. A CMI Controller receives no indication of the interrupt being asserted by the cable assembly's opposite end Free-Side to the Fixed-Side of the CLP. Transferring any indicators from one end to the other is beyond the scope of this specification.

The Free-Side requests interrupt service by continuously asserting the CINT# pin until the Fixed-Side management controller clears the interrupt flag of the Free-Side memory.

The Fixed-Side must provide a pull-up resistor, of value $R_{PU\_AUX}$, on CINT# to MGTPWR to passively de-assert CINT#.

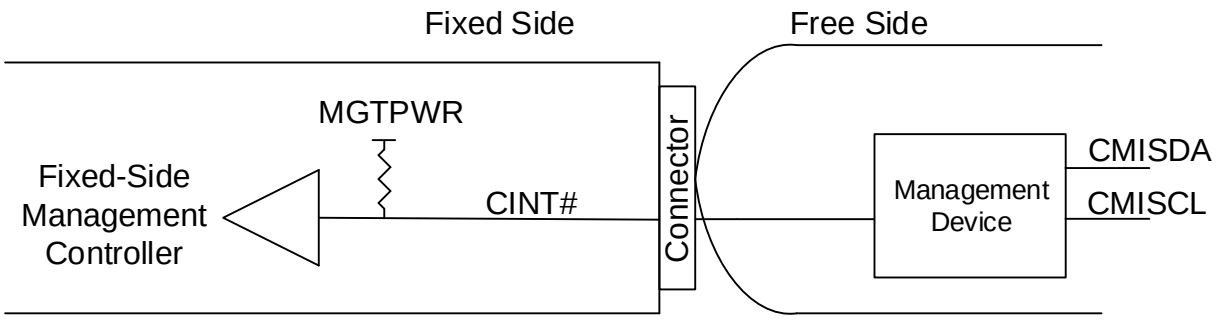


**Figure 2-3. CINT# Signal Implementation**

# 2.6. Cable Management Interface

The Cable Management Interface, or CMI, is a two-wire interface through which various cabled system components communicate with each other. This type of interface is common for Fixed-Side Subsystems to read information from cable assemblies. It is a closed interface meant to only communicate between the cabled Subsystem's management interface components. Multiple devices within a Fixed-side Subsystem must not share CMISCL and CMISDA. A separate interface must be required to interface with the PCI-Express devices and any other system member requiring information from the CMI. Figure 2-4. illustrates an example of a CMI controller implementation of four physical cable ports.

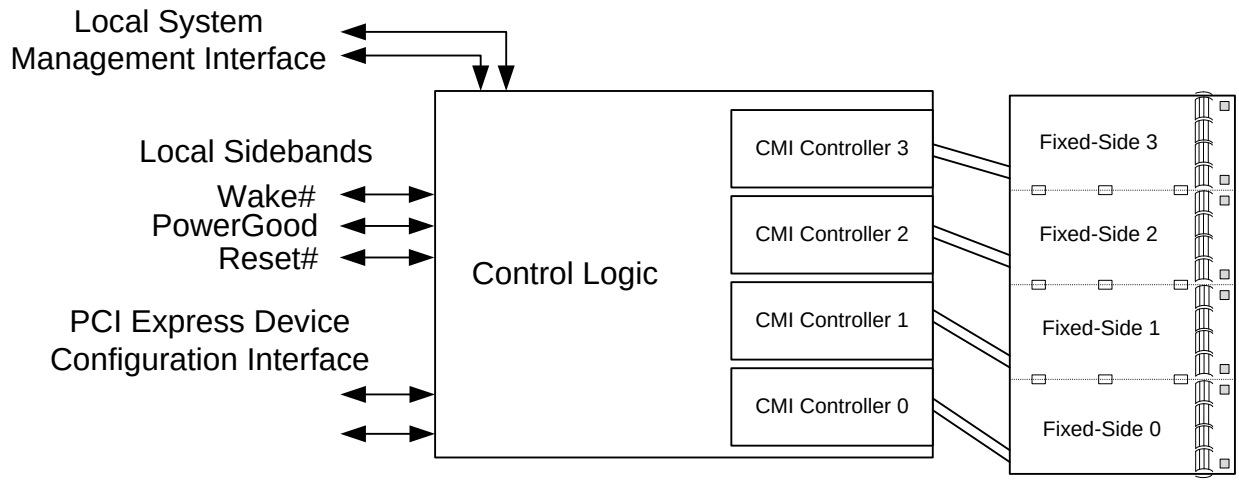


**Figure 2-4. Example CMI Controller for Four x4 Cable Ports**

**Implementation Note:** The CMI may be implemented in programmable logic, micro-controller, or CPU. The specific implementation may be dictated by the features to be supported by the Fixed-Side Subsystem.

Using the Cable Management Interface, each Fixed-Side of the link reads information from the cable assembly to configure the PCI Express interface. Optionally, CMI enables the sending and receiving of sideband messages between the two Fixed-side Subsystems, located at addresses A4h and A6h, if the cable assembly is capable, as indicated by the Sideband Capable bit in the PCIe Capabilities register of the Free-Side. Each Fixed-side Subsystem must implement a multi-master controller to enable communication between the Fixed-Side Subsystems. A list of messages that represent sideband signals is listed in Table 4-2. In addition, vendor-specific messages may also be sent via the Cable Management Interface if both the cable and devices support this operation. To support long cables the bus must operate at 100 kHz as the default data rate.

The CMI must support a clock stretching timeout specified by $T_{CMI\_CLK\_HOLD}$ in Table 2-1.

With the cable management interface, a device in the cable assembly provides manufacturer information, save its state for a suspend event, report certain errors, accept control parameters, and return current status.

The Cable Management Interface conforms to the Management Interface Protocol described by the *PCI-Express OCuLink Specification*. The memory maps are described in section 6.2.1 of this specification.

The Fixed-Side Subsystems provide pull-ups, of value $R_{PU-CMI}$, for CMISCL and CMISDA to the MGTPWR rail, and the components of these signals need to have a 3.3 V signal tolerance. Cable loading limits are specified in Table 2-2. Additional specific parameters for the devices on the CMI interfaces are listed in Table 2-1 and Table 2-2. The falling edge of the waveform must remain below $V_{IL\_CMI}$ after the incident edge until the signal is to transition to a value of one.

**Observation:** Process technologies have allowed significant decreases in the transition time of driver outputs, even for 2-wire interfaces such as CMI. To allow fast edges to properly transition on the falling edge, impedance matching of the CMI device buffer to the $Z_{CMI\_PCB\_ZO}$ impedance target of the fixed side may be required. Some devices may be excluded from this application, including devices with fast edges but high impedance**.**

The memory map for the storage device in the cable assembly is described in section 6.2.2. These tables describe the information stored in the cable assembly that is available to the system. The information needed by the Upstream and Downstream Subsystems for configuration include:

- Cable loss to calculate the starting points for both transmit and receive equalizers
- Bit rates supported by the cable
- Number of Lanes supported by the cable
- Flight time of the cable
- Cable technology

This information is used to configure the fixed-side devices, both logically and electrically. It should be noted that further fine tuning of the transmitters and receivers occurs during the link equalization procedure that is specified in the *PCI Express Base Specification*.

The Power Sequencing section in 4.2.1 describes when the Free-Side memory device is read. The CMI Controllers for each cable port must access the cable assembly memory devices, located at A0h for Downstream devices and A2h or A0h for Upstream devices. Physical cable port 0 is the only port required to be accessed if the Free-Side supports more than a single x4 Link. The Cabled PCIe Capabilities 1, Cabled PCIe Capabilities 2, Cable Technology, and Attenuation registers are used to configure the PCI Express device logical parameters, equalization configuration, and are used for some system parameters. The communication process to configure the PCI Express device and other system parameters, once the information is retrieved and determined, is implementation specific.

For a sideband enabled cable assembly, the Cable Management Interface transmits status change events at the system level. See section 4.1 for a list of status registers that are read and messages that are received.

**Implementation Note:** Not all cable assemblies have bidirectional communication for system status communication via the cable management interface. The implication is that the external device may not be ready when the Host System begins accessing the PCI-Express bus. It is recommended that a user visible indicator be used to tell a user that the external device is ready for the system to begin configuration.

To minimize the impact of topology, loading, and addressing, there is a limit of four loads on the management bus. There is a controller for each fixed-side and two devices in the cable. For both Upstream and Downstream Fixed-Sides, there must be a controller for each individual cable assembly permitted to allow for the creation of a unique link or that may read status of the cable in the lower page of the cable memory device. Downstream

For Upstream or Downstream Fixed-Sides that support either Cable Aggregation or Port Bifurcation, a controller for each cable assembly to be concurrently (or simultaneously) supported (per connector/cable assembly) is required for setting the respective equalizers. If a cable assembly is wide enough for the entire physical link, then only the management data link associated with the lowest lane may be used for all of the management and configuration information. For example, if the port supports sixteen lanes, does not support Port Bifurcation and does not support aggregation, and the cable assembly is monolithic, then only the management controller associated with lanes 0-3 is required. The configuration information from the cable assembly is used to configure all lanes. Figure 2-5 shows the cable management interface connectivity of four independent cables.

**Implementation Note:** If only a single cable assembly is supported by a fixed-side port, configuration data and sideband message are supported for that link, however, the status of active cable components would only be available for the four lowest lanes of the link.

If a multi-connector cable assembly does not implement an interface for quad's greater than zero, then Active Cable Assembly status information for lanes above the lowest quad would not be available to the management controller on the Fixed-Side.

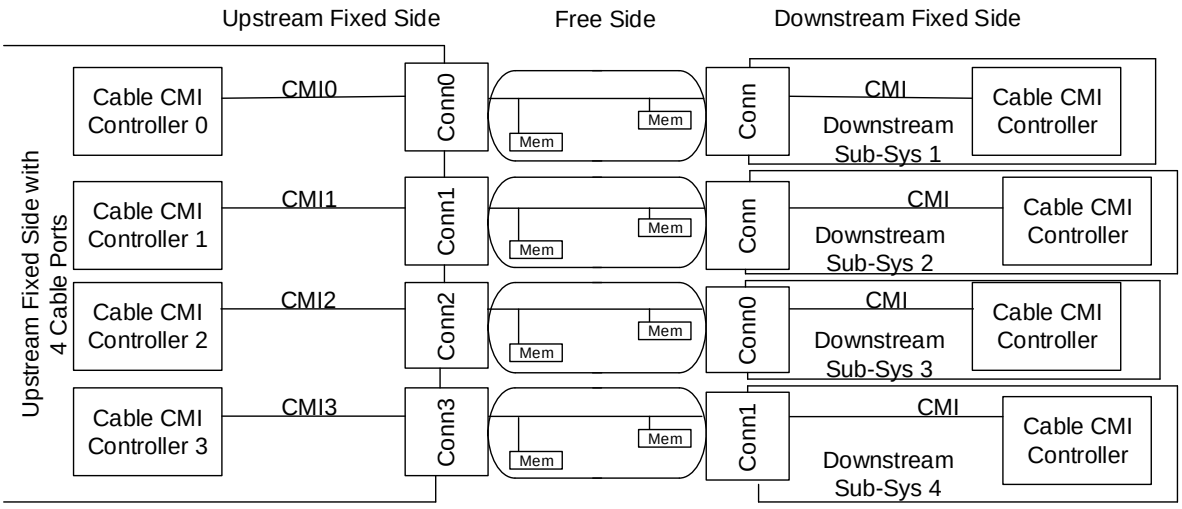


**Figure 2-5. Cable Management Interface Connectivity of Four Independent Cables**

Figure 2-6 illustrates a bifurcated port with two x4 links and one x8 link. The x8 link aggregates two x4 cable assemblies. In this case, there are four Upstream CMI controllers and four Downstream CMI controllers. This x8 Downstream fixed-side needs the CMI controller to support Cable Aggregation and read each cable memory device to access configuration information. However, sideband messages are only required on the CMI link associated with the lowest lane in the link (CMI for Downstream Subsystem 3). The CMI associated with the higher lanes may also send the sideband messages, but each Subsystem must reference the CMI associated with the lowest lane for interpretation of the sidebands for the linked Subsystem.

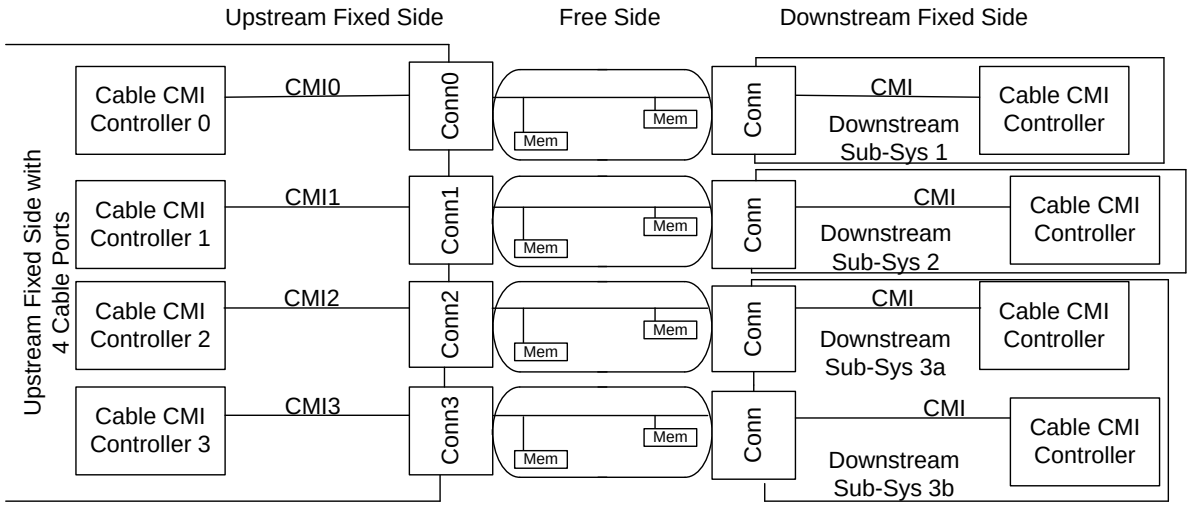


**Figure 2-6. Cable Management Interface Connectivity of Three Independent Cables, with Limited Cable Monitoring Capabilities**

Figure 2-7 is similar, but shows a Downstream x8 fixed-side device that does not support Cable Aggregation and requires the use of x8 cable assembly. The x8 assembly only has a single cable management interface. The Upstream CMI controller must interpret sideband messages received on CMI2 to reference the entire x8 link (the upper two cable connectors). Cable assembly status information is not read from the assembly since neither the cable assembly, nor the Downstream Subsystem, have the components. If the cable assembly contained memory devices, the Upstream port would be able to read cable assembly status information while the Downstream Subsystem would not. The Upstream Subsystem would be able to read the Downstream cable assembly device if the upper quad's management bus traversed the cable assembly.

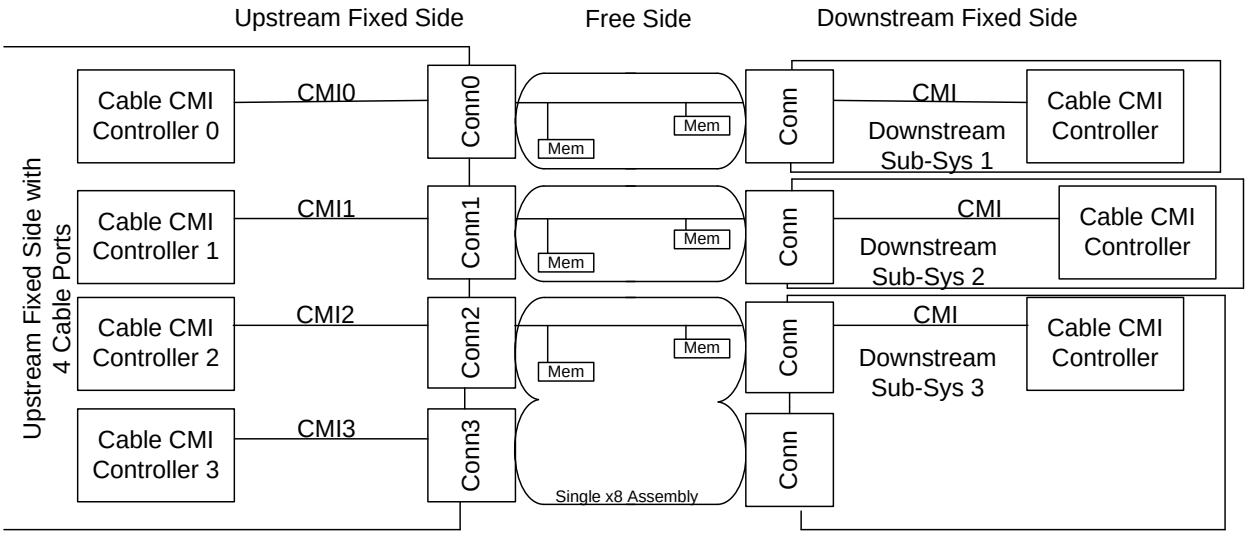


**Figure 2-7. Cable Management Interface Connectivity of three Independent Cables, with limited Cable Monitoring Capabilities.**

## 2.6.1. Capacitive Load of High-power Cable Management Interface Lines

The pin capacitance $C_{OUT}$ is defined as the total capacitive load of one CMI device as seen in a typical manufacturer's data sheet. The capacitive load for a Fixed-Side is defined by $C_{CMI\_FIXED}$. The capacitive load for a cable assembly is defined by $C_{CMI\_CABLE\_LOAD}$.

## 2.6.2. Minimum Current Sinking Requirements for Cable Management Interface Devices

Devices are required to sink a minimum current of $I_{SINK}$ while maintaining the $V_{OL}$ (max) of 0.4 V.

## 2.6.3. Cable Management Interface *Back Powering* Considerations

Unpowered devices connected to either a low-power or high-power CMI segment must provide, either within the device or through the interface circuitry, protection against "back powering" the Cable Management Interface. Unpowered devices must meet $I_{LEAK}$ parameters in Table 2-1.

# 2.7. Electrostatic Discharge

Electrostatic Discharge (ESD) requirements for auxiliary signals are identical to those specified in the PCI Express signal protection. See section 3.5 for ESD requirement details.

# 2.8. Auxiliary Signal Parametric Specifications

Figure 2-8 illustrates the locations where measurements must be taken. Some measurements may require a test fixture at the particular test point for access to the signal and care should be taken to limit or remove the impact of the test fixture.

**FIGURE TBD**

**Figure 2-8. Test Points 1**

## 2.8.1. DC Specifications for Single-ended Signals

Table 2-1 lists the DC specifications for single-ended signals

**Table 2-1. DC Specifications for Single-ended Signals**

| Symbol | Parameter | Min | Max | Units | Notes |
|---|---|---|---|---|---|
| $R_{PU\_AUX}$ | Pull-up Resistor for Fixed-Side Auxiliary signals | 9 k | 10 k | Ω | 2 |
| $R_{PU\_CMI}$ | Pull-up Resistor for Fixed-Side for CMI signals | 2.13 k | 2.17 k | Ω | 3,5,6 |
| $V_{OH\_SB}$ | Voltage output high for cable auxiliary signals | $V_{MGTPWR}$-0.5 | $V_{MGTPWR}$+0.3 | V | 1,2 |
| $V_{OL\_SB}$ | Voltage output low for cable auxiliary signals | 0 | 0.4 | V | 1,2 |
| $V_{IH\_SB}$ | Voltage input high for cable auxiliary signals | $V_{MGTPWR}$ * 0.7 | $V_{MGTPWR}$+0.5 | V | 1,2 |
| $V_{IL\_SB}$ | Voltage input low for cable auxiliary signals | -0.3 | $V_{MGTPWR}$ * 0.3 | V | 1,2 |
| $V_{OH\_CMI}$ | Voltage output high for CMI signals | $V_{MGTPWR}$-0.5 | $V_{MGTPWR}$+0.3 | V | 1,3 |
| $V_{OL\_CMI}$ | Voltage output low for CMI signals | 0 | 0.4 | V | 1,3,6 |
| $V_{IH\_CMI}$ | Voltage input high for CMI signals | $V_{MGTPWR}$ * 0.7 | $V_{MGTPWR}$+0.5 | V | 1,3 |

| Symbol | Parameter | Min | Max | Units | Notes |
|---|---|---|---|---|---|
| $V_{IL\_CMI}$ | Voltage input low for CMI signals | -0.3 | $V_{MGTPWR}$ * 0.3 | V | 1,3,6 |
| $I_{SINK}$ | Current a CMI device must be able to sink | 3.0 | | mA | |
| $I_{LEAK}$ | Input Leakage Current | | ±5 | µA | 3 |

**Notes:**

1. Measured at the Fixed-Side Subsystem connector.
2. Applies to CBLPRSNT#, CINT# and CADDR
3. Applies to CMISDA and CMISCL.
4. Capacitance includes devices and routing path.
5. Pull-up resistors are implemented by both Fixed-Sides, so two pull-up resistors in parallel are used to generate the high levels and the drivers must be able to meet low voltages with the lower equivalent resistance.
6. The value needed to balance cable capacitive limits and timings require a 2.15 k 1% resistor tolerance.

## 2.8.2. AC Specifications for Single-ended Signals

Refer to Appendix B in the *PCI Express OCuLink Specification* Revision 1.0 for most of the CMI timing requirements. Table 2-2 lists additional requirements to those listed in the *OCuLink Specification*.

## Table 2-2. AC Specifications for Single-Ended Sideband Signals

| Symbol | Parameter | Value | Units | Notes |
|---|---|---|---|---|
| $F_{CMI}$ | CMI Operational Frequency | 100 | kHz | 1 |
| $T_{CMI\_CLK\_HOLD}$ | Serial Interface Clock Hold-off (Clock Stretching) | <500 | µs | 2,3 |
| $T_{MGT\_INIT}$ | Management Interface Initialization Time | <500 | ms | 4 |
| $T_{SB\_TIMEOUT}$ | Time limit between sending CMI message update | <500 | ms | |
| $Z_{CMI\_PCB\_ZO}$ | Impedance of CMI PCB traces | 55 | Ω | 5 |
| $Z_{CMI\_CABLE\_ZO}$ | Impedance of CMI Wires in Cable Assembly | >50 | Ω | 6 |
| $C_{CMI\_CABLE\_LOAD}$ | Total capacitance of each CMI signal in a Cable Assembly | <330 | pF | 7 |
| $T_{CMI\_SUPRESS}$ | Width of noise spikes which an input filter must suppress on the CMI inputs | 50 | V | |
| $V_{HYSTERESIS}$ | Hysteresis of Schmitt trigger inputs | .05* $V_{MGTPWR}$ | V | |
| $C_{CMI-FIXED}$ | Capacitance of CMI Interface for the Fixed-Side | 30 | pF | 5,8 |

**Notes:**

1. See section 2.6 for information on the behavior of $F_{CMI}$. This excludes clock stretching.
2. Amount of time the cable or CLP is permitted to hold CMISCL low before continuing with a read or write operation.
3. It is recommended that clock stretching by the controller be disabled if the Free-Side is optical. See 4.2.1.2.
4. SFF-8449 cables specify a $T_{MGT\_INIT}$ of 2000 ms.
5. Applies to CMISCL and CMISDA traces on Fixed-Side PCB.
6. Applies to CMISCL and CMISDA wires in the cable assembly.
7. Total capacitive load of each cable assembly's CMI signals, including components.
8. Pull-up resistors are implemented by both Fixed-Sides, so two pull-up resistors in parallel are used to generate the high levels and the drivers must be able to meet low voltages with the lower equivalent resistance.

# 2.9. Power Requirements

Power supply filtering per cable is recommended to minimize system noise from impacting the circuits in the cable assembly, or for one cable assembly impacting another. Table 2-3 lists the PWR and MGTPWR requirements.

## Table 2-3. PWR and MGTPWR Requirements

| Symbol | Parameter | Value | Units | Notes |
|---|---|---|---|---|
| $V_{MGTPWR}$ | Management Power Pin Voltage | 3.0-3.6 | V | 4 |
| | Maximum Sustained Current | 30 | mA | |
| $V_{PWR}$ | Power Pin Voltage | 3.3 ± 5% | V | 1, 4 |
| | Power supply noise tolerance (Peak) | <50 | mV | 4 |
| | Power Consumption for Free-Side | 1.5 | W | 3, 5 |
| | Maximum Sustained Current per $V_{PWR}$ Pin | <500 | mA | 1, 2 |
| | Capacitive Load of the Cable Assembly per $V_{PWR}$ Pin | <15 | µF | |

**Notes:**

1. This is per pin.
2. Current requirements apply to the current through the contacts.
3. Maximum power consumption of the cable assembly must not exceed 1.5 W from 500 ms after power up.
4. Includes noise and ripple
5. This is the total power limit for the module. How power is consumed per pin is beyond the scope of this specification.

An example of a filter implementation for the power pins per connector plug receptacle power pin is shown in Figure 2-9. It is not a recommended implementation and provided for informational purposes only. For this example, the filter inductor is 1 µH with an ESR of 100 mΩ. The bulk capacitors are 22 µF with an ESR of 220 mΩ.

It is strongly recommended that each power pin be filtered independently.

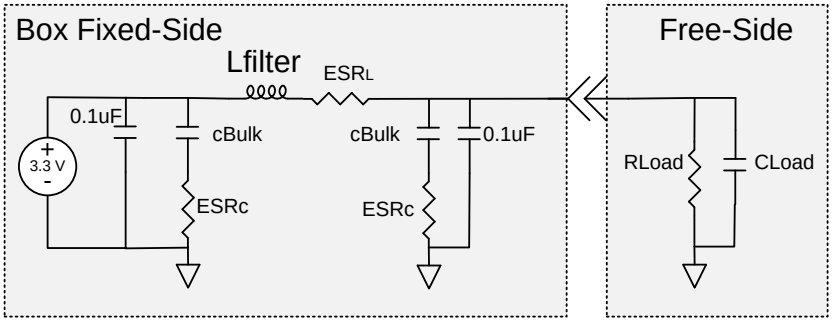


## Figure 2-9. Filter Implementation Example

589

3

590 # 3. PCI Express Signals

591 Each PCI Express Lane consists of a pair of differential signals. The Transmitter pair is labeled
592 PETpN and PETnN where **N** is the lane number (starting with 0); **p** is the true signal and **n** is the
593 complement signal. The Transmitter pair originates in the Upstream Subsystem and is connected to
594 a Receiver pair at the Downstream Subsystem at the opposite side of the cable. The Receiver pair is
595 labeled PERpN and PERnN with the same labeling convention as the Transmitter pair.

596 Support for Polarity Inversion is required on all PCI Express Receivers across all Lanes
597 independently. The **p** and **n** connections may be reversed to simplify trace routing and minimize
598 vias if needed. All PCI Express Receivers incorporate automatic Polarity Inversion as part of the
599 Link Initialization and Training and corrects the polarity independently on each differential pair.
600 Refer to Chapter 4 of the *PCI Express Base Specification* for more information.

601 Support for Lane Reversal is optional. An example of Lane Reversal Lane 0 of an Upstream Port
602 connected to Lane N-1 of a Downstream Port where either the Downstream or Upstream device
603 may reverse the Lane order to configure an **xN** Link. Refer to Chapter 4 of the *PCI Express Base*
604 *Specification* for more information.

605 A *null* modem function for connecting the transmit pair on one side to the receive pair to the other
606 side is performed within the cable assembly. External cable connector pin assignments are identical
607 at both sides of the Link. Lane Polarity Inversion and/or Lane reversal must not be implemented
608 within the cable assembly.

609 ## 3.1. Interconnect

610 In the context of this specification, *interconnect* comprises everything between the pins of the
611 Transmitter package and the pins of the Receiver package. This consists of traces on printed circuit
612 boards, cable, AC-coupling capacitors, and connectors. The interconnect total capacitance to ground
613 seen by the Receiver Detection circuit (see Chapter 4 of the *PCI Express Base Specification*) must not
614 exceed 3 nF, including capacitance added by attached test instrumentation. Note that this
615 capacitance is separate and distinct from the AC-coupling capacitance value (see section 3.2.1).

## 3.1.1. Link Definition

Typical cabled PCI Express Links, from source to destination, consist of the following:

- Transmitters on an ASIC on a printed circuit board
- Package fan-in-out trace topologies
- PCB coupled microstrip and/or stripline traces
- AC-coupling capacitors
- Vias for layer changes
- Cable mated connector
- External raw cable
- Coupled microstrip line and/or stripline traces
- Receivers on an ASIC on a printed circuit board

Figure 3-1 shows the electrical parameters for the Link are subdivided into:

- Upstream Subsystem
- Cable including mated connectors
- Downstream Subsystem

Note that for validation purposes, the separation is somewhat different to facilitate ease of connecting any test and measurement equipment.

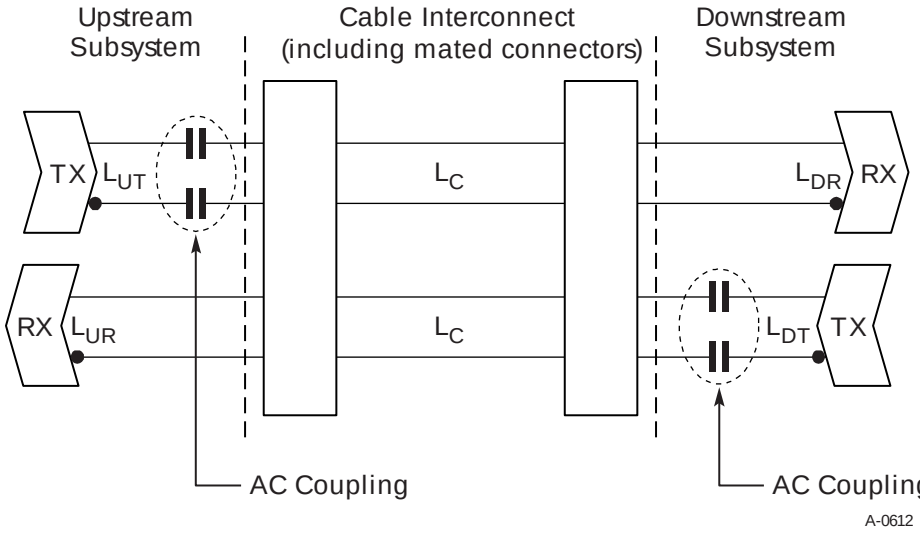


**Figure 3-1. Electrical Parameter Allocation**

# 3.2. Electrical Budget

## 3.2.1. AC-coupling

AC-coupling is required at the Transmitter Fixed-Side (see Figure 3-1) with values as specified in Chapter 4 of the *PCI Express Base Specification*. A passive cable assembly does not implement any AC-coupling capacitors. Active cables must provide AC-coupling capacitors in an implementation specific manor.

**Note:** SATA and SAS systems implement the capacitors on the Rx Subsystem and not the Tx Subsystem. With both specifications utilizing the SFF-8644 connector, it is possible to cross-plug a PCI Express port with a SAS port. It is recommended that PCI Express receivers tolerate a DC connection to SATA or SAS transmitters without damage. The method used to achieve this is beyond the scope of this specification

## 3.2.2. Jitter from Clock Sources

A portion of the link budget is required for the differences in the clocks of the two fixed Subsystems. The jitter requirements for the Upstream and Downstream clocks that are referenced by the cable link are in Table 3-1.

### Table 3-1. Jitter Budgeting Assumptions for Separate Clocks

| Jitter Parameters | Symbol | Values (RMS in ps) | Notes |
|---|---|---|---|
| Link Jitter for 2.5 GT/s | $T_{INDCLK\_LINK\_JITTER\_2.5G}$ | 4.30 | |
| Link Jitter for 5.0 GT/s | $T_{INDCLK\_LINK\_JITTER\_5G}$ | 2.0 | Likely to be removed |
| Link Jitter for 8.0 GT/s | $T_{INDCLK\_LINK\_JITTER\_8G}$ | 0.50 | |

**Notes:**
1. Defined for either SSC or non-SSC implementations.
2. Filter functions are defined in the *PCI Express Base Specification*.

## 3.2.3. Crosstalk

All Subsystem designs must properly account for any crosstalk that may exist among the various differential signal pairs and other signals alike. Crosstalk is either near-end (NEXT) or far-end (FEXT). Each crosstalk component has potential impact on a design and must be planned for accordingly. Jitter budgets assigned to the Subsystems are set at an absolute minimum to preserve the maximum possible budget for the copper cable interconnect.

Crosstalk between differential pairs within the interconnect path influence and impact the data signals and any subsequent loss and jitter. The eye diagrams in section 3.3 account for any and all crosstalk allowed.

## 3.2.4. Lane-to-Lane Skew

The skew at any point is measured at the zero crossings of differential voltage of the compliance pattern, while simultaneously transmitting on all Lanes (see Table 3-2). Refer to the *PCI Express Base Specification* for a definition of the compliance pattern.

### Table 3-2. Allowable Interconnect Lane-to-Lane Skew

| Jitter Parameters | Symbol | Values (RMS in ps) | Notes |
|---|---|---|---|
| Total Interconnect Skew | $S_T$ | 2.7 ns | 1, 2, 3, 6 |
| Subsystem Skew | $S_S$ | 0.35 ns | 4 |
| Cable Assembly Skew | $S_C$ | 2.0 ns | 5 |

**Notes:**

1. This does not include Transmitter output skew.
2. This is the total skew allowed for a link. If multiple cable assemblies are used, the flight time delay information (Table 6-3 bytes 108-109) must be read from all aggregated cable assemblies in the link and must meet the Total Interconnect Skew for the entire link.
3. The skew at any point is measured at the zero crossings of differential voltage of the compliance pattern, while simultaneously transmitting on all lanes. The compliance pattern is defined by the PCI Express Base Specification.
4. The limit is based on approximately two inches of stripline.
5. This limit is across all signal pairs. When using multiple cable assemblies for a port aggregation, the total skew across all pairs in the link must meet this requirement.

## 3.2.5. Transmitter De-Emphasis

De-emphasis is required in the Transmitter for 2.5 GT/s and 5.0 GT/s to reduce ISI. As stated in Chapter 4 in the *PCI Express Base Specification*; for full swing mode this is implemented as a -3.5 dB (±0.5 dB) attenuation of all non-transition bits relative to the amplitude of the preceding transition bit when operating at 2.5 GT/s. For 5.0 GT/s, the *PCI Express Base Specification*, Revision 2.0 added an additional setting of -6.0 dB (±0.5 dB) for full swing mode. For cabled implementation, transmitters must be configured for the -6.0 dB de-emphasis setting. While not optimal for all cable lengths, sufficient margin exists when using low-loss cables for reliable operation with the -6.0 dB setting.

For 8.0 GT/s, a more sophisticated transmitter equalization scheme is required. The *PCI Express Base Specification*, Revision 3.0 describes the equalization requirements, which includes not only more complex transmitter equalization, but adds receiver equalization as well. This specification builds upon these characteristics to implement this cabled form factor.

### 3.2.6. Skew Within the Differential Pair (Intra-Pair Skew)

Skew within a differential pair (intra-pair skew) gives rise to a common-mode signal component, which increases Electromagnetic Interference (EMI). The differential pair(s) on a Subsystem printed circuit board should be routed such that the skew within each differential pair is ≤0.005 inch.

Intra-pair skew of the cable assembly is more difficult to control and tight specifications result in increased cost. A maximum skew of 0.2 UI is recommended for the cable assembly intended to operate at 2.5 GT/s. It is left up to the application to make appropriate cost/performance tradeoffs. No specific budgets are provided for intra-pair skew of the cable assembly as these are incorporated within the overall cable assembly budgets.

For 5.0 GT/s and faster, the skew is not explicitly defined for the cable. Instead, a limit is placed on mode conversion for the cable assembly in section 6.2.4.

## 3.3. Eye Diagrams

The eye diagrams defined in this section represent the compliance eye diagrams that must be met for both Upstream and Downstream Subsystems. Transceiver silicon requirements are as specified in the *PCI Express Base Specification*.

### 3.3.1. Fixed-Port Transmit Compliance Eye

The Transmitter compliance eye is defined by the values in Table 3-3 for 2.5 GT/s and Table 3-4 for 8.0 GT/s. Figure 3-2 shows transmitter compliance eye diagram

**Table 3-3. Transmitter Path Compliance Eye Requirements at 2.5 GT/s**

| Parameters | Values | Notes |
|---|---|---|
| $Vtx_A$ | TBD | 1, 2, 3, 4, 5, 6, 7 |
| $Vtx_A$_d | TBD | 1, 2, 3, 4, 5, 6, 7 |
| $Ttx_A$ @ BER $10^{-12}$ | TBD | 1, 2, 3, 4, 5 |
| $Ttx_A$ @ $10^6$ Samples | TBD | 1, 2, 3, 4, 5 |

**Notes:**

1. The signal acquired from TP2.
2. An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating this eye diagram (Figure 3-2).
3. Transition and non-Transition bits must be distinguished to measure compliance against the de-emphasized voltage level ($Vtx_A$_d).
4. $Ttx_A$ is the eye width.
5. The values in this Table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median.
6. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.
7. Transition bits measured with transmitter set to -3.5 dB of de-emphasis. Transition and non-Transition bits must be distinguished to measure compliance against the de-emphasized voltage level ($Vtx_A$_d).

695 **Table 3-4. Transmitter Path Compliance Eye Requirements at 8.0 GT/s**

| Parameters | Values | Notes |
|---|---|---|
| $Vtx_A$ | >34 mV | 1, 2, 4, 5 |
| $Vtx_A$_d | TBD | 1, 2, 4, 5 |
| $Ttx_A$ @ BER $10^{-12}$ | >41.2 ps | 1, 2, 3, 4 |
| $Ttx_A$ @ $10^6$ Samples | TBD | 1, 2, 3, 4 |

**Notes:**

1. An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating this eye diagram (Figure 3-2).
2. Transition and non-Transition bits must be distinguished to measure compliance against the de-emphasized voltage level ($Vtx_A$_d).
3. $Ttx_A$ is the eye width.
4. The values in this Table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median.
5. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.

696

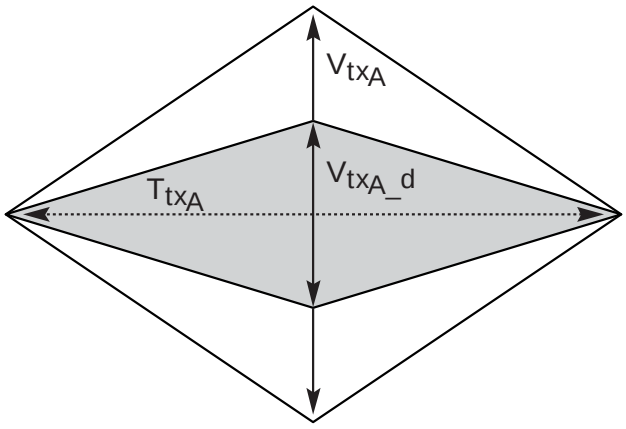


697 A-0615

698 **Figure 3-2. Transmitter Compliance Eye**

700 # 3.3.2. Cable Port Transmit Compliance Eye

701 The minimum sensitivity for the Subsystem's Receiver path compliance is specified in Table 3 5 for
702 2.5 GT/s and Table 3 6 for 8.0 GT/s. A representative eye diagram is shown in Figure 3 3.

703 **Table 3-5. Receiver Path Compliance Eye Requirements at 2.5 GT/s**

| Parameters | Values | Notes |
|---|---|---|
| $Vrx_A$ | TBD | 1, 2, 3, 4, 6, 7 |
| $Vrx_A$_d | TBD | 1, 2, 3, 4, 6, 7 |
| $Trx_A$ @ BER $10^{-12}$ | TBD | 1, 2, 3, 4, 5 |
| $Trx_A$ @ $10^6$ Samples | TBD | 1, 2, 3, 4, 5 |

**Notes:**

1. Requirements for the signal acquired from TP4.
2. All signal and timing values are referenced at the Subsystem connector mounting pads.
3. An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating the eye diagram shown in Figure 3-3.
4. Transition and non-Transition bits must be distinguished to measure compliance against the de-emphasized voltage level ($Vrx_A$_d).
5. $Trx_A$ is the eye width.
6. The values in this Table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median.
7. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.

704

705 **Table 3-6. Receiver Path Compliance Eye Requirements at 8.0 GT/s**

| Parameters | Values | Notes |
|---|---|---|
| $Vrx_A$ | > 34 mV | 1, 2, 3, 4, 6, 7 |
| $Vrx_A$_d | TBD | 1, 2, 3, 4, 6, 7 |
| $Trx_A$ @ BER $10^{-12}$ | > 41.2 ps | 1, 2, 3, 4, 5 |
| $Trx_A$ @ $10^6$ Samples | TBD | 1, 2, 3, 4, 5 |

**Notes:**

1. Requirements for the signal acquired from TP4.
2. All signal and timing values are referenced at the Subsystem connector mounting pads.
3. An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating the eye diagram shown in Figure 3-3.
4. Transition and non-Transition bits must be distinguished in order to measure compliance against the de-emphasized voltage level ($Vrx_A$_d).
5. $Trx_A$ is the eye width.
6. The values in this Table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median.
7. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.

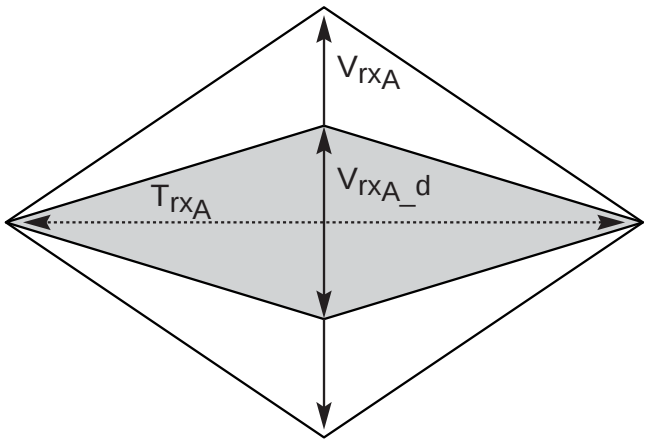


A-0616

**Figure 3-3. Receiver Compliance Eye**

# 3.4. Receiver Path Sensitivity Testing

The minimum sensitivity values for the system board Receiver path compliance at 8.0 GT/s are defined in Table 3-7 and Table 3-8. The receiver path must be tested with a worst-case eye in order to verify that it achieves a BER $<10^{-12}$. This worst-case eye is calibrated using TX equalization settings that are optimal with the reference equalizer for the calibration channel. After calibration, the test-generator's equalization settings may be adjusted using the transmitter equalization setting in the required TX equalization space preferred by the device under test, without changing any other parameter of the test signal or recalibrating the test signal.

**Note:** If the test generator's TX equalization settings are adjusted away from the optimal settings and the test generator is not able to change transmitter equalization without impacting other calibrated parameters, then the other parameters must be adjusted back to the specified values.

If the test is not run in a way that produces the worst case cross-talk that would be present with all lanes active, the additional cross-talk must be accounted for in some other way. While the receiver's capacity to adapt its own equalization is part of the test described above, its ability to request the link partner's transmitter to change its equalization settings is tested by applying a signal whose equalization settings are sub-optimal compared to the jitter sensitivity test signal described above. For this signal, the reference receiver would not be able to achieve proper equalization by means of its own CTLE and DFE alone. Such a signal is defined using the signal resulting from the calibration method described above and adjusting the test-generator equalization. Note that if the RX under test is more capable than the reference (CTLE+DFE) receiver, the RX may not require the TX to change its equalization levels to achieve a BER $< 10^{-12}$. In any case, equalization settings resulting from this procedure must be used for the above RX test and, if the RX requires the TX equalization to change, such a change must be accommodated by the test set-up used. A specific methodology for this procedure is outside the scope of this specification.

718 **Table 3-7. Free-Port Receiver Path Sensitivity Requirements at 8.0 GT/s**

| Parameter | Min | Max | Unit | Comments |
|---|---|---|---|---|
| VRX-EH-8G Eye height | > 34 | | mV | |
| TRX-EH-8G Eye width | > .33 | | UI | |
| Rj (Random Jitter) | 3 | | Ps RMS | |
| Sj (Sinusoidal Jitter) 100 MHz | 12.5 | | Ps | |
| Differential Mode Sinusoidal Interference 2.1 GHz | 14 | | Ps | |

**Note:** These are the requirements for the signal inserted into TP1.

719

720 **Table 3-8. Fixed-port Receiver Path Sensitivity Requirements at 8.0 GT/s**

| Parameter | Min | Max | Unit | Comments |
|---|---|---|---|---|
| VRX-EH-8G Eye height | > 34 | | mV | |
| TRX-EH-8G Eye width | > .33 | | UI | |
| Rj (Random Jitter) | 3 | | Ps RMS | |
| Sj (Sinusoidal Jitter) 100 MHz | 12.5 | | ps | |
| Differential Mode Sinusoidal Interference 2.1 GHz | 14 | | ps | |

**Note:** These are the requirements for the signal inserted into TP3.

721 # 3.5. ESD

722 PCI-Express cable ports of all link widths must withstand 2 kV of ESD contact discharge to the
723 connector cage using the human body model (HBM), Class 2 per *JEDEC JESD 22-A114:B 2000*
724 with power applied, without damage not limited to latch up, without a cable inserted, and without
725 non-recoverable errors with a cable inserted.

726 A recoverable error is one that does not require reset or replacement of the device.

727

728

4

729 # 4. Interoperability

730 Expanding a system utilizing PCI-Express over a cable requires some understanding of the
731 implications to interoperability. Solutions with cables that are external to the Root Complex
732 Subsystem's software or power control present some unique challenges to the PCI Express
733 environment. A new management interface has been developed to comprehend these challenges.
734 Flexibility in the cable assembly usage model allows users some choice in cable features and the
735 ability to interface with Downstream Subsystems that conform to the *PCI-Express External Cabling*
736 *Specification*, Revision 2.x or prior.

737 The Cable Management Interface across the cable handles sideband communication and enables
738 power sequencing, configuration, reading of status, and an optional communications link in tandem
739 to the PCI-Express link. It is an optional feature that allows the use of cables that adhere to SFF-
740 8449 specification to be used in an externally cabled PCIe environment with a reduced feature set.
741 Active Optical Cable assemblies may not implement CMI across the cable for cost or complexity
742 reasons, and therefore optical cables may have a reduced feature set. The Upstream Subsystems
743 should be designed in such a way as to not require the use of CMI across the cable. However, the
744 CMI controller is still required by both the Upstream and Downstream fixed ends to read cable
745 assembly information for configuration of the PCIe devices that are part of the cabled link.
746 Controlled power sequencing is only achieved with sideband enabled cable assemblies. Some
747 features are only enabled using the sideband interface, such as SRIS or remote power control.

748 The Legacy adapter cable has a connector with a memory device for the Upstream end and a Legacy
749 connector for the Downstream end. This means the cable assembly is directional. Since this version
750 of the specification does not have all the same sideband mechanisms that the previous revisions
751 contained, several of these sidebands must be recreated within the cable assembly.

752 This usage of Extension devices is beyond the scope of this specification.
753

**Note:** Extension devices are limited to two per link. Extension device usage is not bounded by form factor specifications. This allows extension devices to bridge form factors, such as an expansion card interface to a cabled interface and a potential unknown number of extension devices in a link.

754

# 4.1. Cable Management Interface

Sideband signals are sent to the link partner using the Cable Management Interface. After the initial link configuration information has been obtained from the cable assembly, sideband messages may be sent from one link partner to the other if the cable supports it. A message should be sent anytime a change occurs at the local Fixed-Side. The intent is to indicate a change in status or that an action needs to be taken by the far end Fixed-Side. Sideband messages only need to be sent and monitored over the CMI for the lowest lane connector in a link that is greater than x4. The lowest lane is defined by the pin assignment in this specification and is not affected by lane reversal.

The data contained in the Vendor Message Space is beyond the scope of this specification.

The Upstream fixed device address target is A6h.

The Downstream fixed device address target is A4h.

## 4.1.1. Fixed-side Memory Map

The fixed Subsystems have memory to provide information for configuration, status, and Vendor-Specific Messages. The Read or Write status is relative to the cable link and does not describe how it is programmed or read by the CMI controller. Table 4-1 lists the Fixed-side memory map values.

### Table 4-1. Fixed-side Memory Map

| Offset | Description | Values | R/W | Notes |
|---|---|---|---|---|
| 0-1 | Sideband Message Register | Bit 0 – CMI_RESET<br>Bit 1 – CMI_POWER_ENABLE<br>Bit 2 – CMI_SSTART<br>Bit 3 – CMI_CLP_READY<br>Bit 4 – CMI_CLP_PRESENT<br>Bit 5 – CMI_WAKE<br>Bit 6 – CMI_HOT-PLUG ATTN BUTTON DETECT<br>Bit 7 – CMI_HOT-PLUG INDICATOR REQ<br>Bits 8:15 – Reserved | W | The Cabled Link Partner management controller writes sideband messages to these bits to indicate a status change of a sideband or indicate an action |
| 2-3 | Cabled Port Status Register | Bit 0 – CPSR_RESET<br>Bits 1:2 – Reserved<br>Bit 3 – CPSR_CLP_READY<br>Bit 4 – CPSR_CLP_PRESENT<br>Bits 5:15 – Reserved | R | These bits are controlled by the Fixed-Side to indicate its status. The opposite fixed-side controller may read these bits to obtain status. |
| 4-5 | PCIe Cable Port Capabilities Register | Bit 0 – 5.0 GT/s Data Rate Supported<br>Bit 1 – 8.0 GT/s Data Rate Supported<br>Bits 2:4 – Reserved.<br>Bit 5 – SRIS Support<br>Bits 6:15 – Reserved | R | These bits are controlled by the Fixed-Side and used by the opposite side during configuration. Byte 4 Bit 1 should always be set to 1. |
| 6-7 | Vendor ID | PCI SIG defined Vendor ID in hex | R | Use for identifying common boxes for port configuration. Recommended to match SVID. |
| 8-9 | Device ID | Vendor defined Device ID in hex | R | Use for identifying common boxes for port configuration. Recommended to match SSID. |

| Offset | Description | Values | R/W | Notes |
|---|---|---|---|---|
| 10 | Connector Configuration | Bits 0:1 – Physical port ordering of x4 connectors, with 00b being the location of Lanes 0-3 and 11b containing Lanes 12-15 of a x16 link.<br>Bits 2:7 – used to identify connector ports that are permitted to be combined into a single link. | R | Lower two bits indicate physical port ordering of x4 ports that may be used to create a larger link, up to x16 lanes or 4 physical ports.<br>Upper six bits indicate physical ports that may belong to a single link. For Upstream Fixed-side devices, physical ports are allowed to bifurcate into a maximum of four individual link connections. |
| 11-18 | Serial Number | Vendor defined 64-bit field to uniquely identify each unit of a model. | | Uniquely identify vendor boxes of same VendorID / DeviceID. Used for port configuration connectivity (mis-wiring). |
| 19-42 | Vendor ID Text | | R | Vendor-defined 24 Character ASCII Value. |
| 43-66 | Vendor Description Text | | R | Vendor-defined 24 Character ASCII Value. |
| 67-126 | Reserved | | R | |
| 127 | CRC | | R | TBD |
| 128-255 | Vendor Message Space | Vendor Specific | W/R | This space is used for vendor-specific communication |

771

**Implementation Note:** This data in the Vendor Message Space may include pointers to additional memory that a vendor may wish to take advantage of. Usage of this space is beyond the scope of this specification.

772

## 773 4.1.2. Sideband Messages

774 Sideband messages are used as an economical alternative to connector pins and cable wires. This
775 method enables the transmission and reception of sideband messages to update a status, indicate an
776 action, or inform the CLP of a local status change. A message must be sent by the CMI controller
777 when a status change occurs in the Fixed-side Subsystem that is to be reflected via a sideband. If
778 transmission of the message fails, the CMI controller must immediately attempt to resend the
779 sideband message until the transmission is successful.

780 When sideband messages are sent by the CMI controller, any bits not specified to change must
781 reflect their state from either their last transmission or initialization.

782 A subsequent message, with no changes from the previous message, must be sent within $T_{SB\_TIMEOUT}$
783 of the previous message.

784 **Table 4-2. Sideband Messages**

| Sideband Message Bits | Bit | Values | R/W | Notes |
|---|---|---|---|---|
| CMI_RESET | 0 | Set indicates PCIe Reset is asserted<br>Cleared indicates PCIe Reset is de-asserted | W | This signal indicates the state of the reset in the root complex box. This is sent by the Upstream Fixed-side only. Sent by Downstream Fixed-side is undefined. |
| CMI_POWER_ENABLE | 1 | Set indicates that main power should be applied,<br>Cleared indicates main power should be off. | W | Indicates the directed power state. This is sent by the Upstream Fixed-Side only. Sent by Downstream Fixed-Side is undefined. |
| CMI_SSTART | 2 | Set indicates a box should wait to enable main power until a sideband message is received with CMI_CLP_READY set.<br>Cleared indicates that main power may be enabled as soon as directed. | W | This bit allowed staggered power up of Boxes that are logically connected via PCI-Express, but do not share a common power delivery or control structure. This is sent by the Upstream Fixed-side only. Sent by Downstream Fixed-Side is undefined. |
| CMI_CLP_READY | 3 | Set indicates that Cable Link Partner has main power enabled, clock is stable. | W | The source is ready for communication that power is good, clock is stable. This is a status signal and does not indicate action. This may be sent by either end of the cabled link. |
| CMI_CLP_PRESENT | 4 | | W | This indicates to the CLP that the local box is present. |
| CMI_WAKE | 5 | Set to request that upstream Fixed side enables main power. | W | This is driven by the Downstream Fixed-Side only. Sent by Upstream device is undefined. |
| CMI_HOT_PLUG_ ATTENTION_BUTTON | 6 | Cleared- No event to report<br>Set - Hot-Plug Request has been issued | W | This bit indicates that the opposite fixed-side has detected a button press. |
| CMI_HOT_PLUG_ INDICATOR_REQ | 7 | Cleared: Normal Operation<br>Set: Hot-Plug Indicator Request | W | This bit indicates that the opposite fixed-side has received direction to blink the indicator. |
| Reserved | 8-15 | | | |

**Note:** Due to the slow nature of the Cable Management Interface, this is not able to effectively implement OBFF via the sideband messages.

785 ## 4.1.3. CMI_RESET

786 The CMI_RESET bit indicates the status of the PCI-Express reset of the Upstream Fixed-side
787 Subsystem. The bit is Set when local reset is asserted, and Cleared when de-asserted. This is sent by
788 the Upstream Fixed-side only. All DSF CMI controllers must send messages with CMI_RESET Set
789 from controller initialization until a message is received on the USF port with it cleared. When sent
790 by the Downstream Fixed-side Subsystem the meaning is undefined.

### 4.1.4. CMI_POWER_ENABLE

The CMI_POWER_ENABLE bit controls the power-up of Boxes that are logically connected via PCI-Express but do not share a common power delivery or control structure. This is sent by the Upstream Fixed-side only. When sent by Downstream Fixed-Side Subsystem is undefined.

**Implementation Note:** Boxes with only DSF ports send this as a reflection of a local POWER_GOOD signal. However, if this simply reflects POWER_GOOD then a predetermined power up sequence is not enabled and one or more boxes in the cabled hierarchy may not be ready to respond to a request with a Successful Completion status.

### 4.1.5. CMI_SSTART

The CMI_SSTART bit in the message is to be Set in conjunction with either CMI_POWER_ENABLE or CMI_WAKE being Set. If sent in a message without a transition of either CMI_POWER_ENABLE or CMI_WAKE being Set, the meaning is undefined. The bit may be Cleared any time after a message has been received with CMI_RESET de-asserted.

### 4.1.6. CMI_CLP_READY

CMI_CLP_READY is sent by the CMI controller to the CLP to indicate the local Box is ready to receive a message with CMI_RESET de-asserted. This implies that local power is stable and within operating parameters, clocks are stable, the configuration data from the cable assembly has been obtained, a message with CMI_CLP_READY has been received on all DSF ports that are capable of receiving sideband messages, and any other implementation-specific criteria has been met for the Box to begin receiving PCI Express transactions. DSF ports that are capable of receiving sideband messages are defined by having CBLPRESNT# asserted and the inserted cable is sideband-capable.

The CMI_CLP_READY signal must remain Set until the unit is no longer able to accept transactions on the PCI Express interface. It is permitted to Clear CMI_CLP_READY by the CMI controller for implementation specific reasons, such as loss of key functionality. However, other Boxes do not have a specified behavior upon receipt of a message with CMI_CLP_READY Cleared after the power up sequencing.

### 4.1.7. CMI_WAKE

This is driven by the Downstream Fixed-Side System only. This message is undefined when sent by an Upstream device. This bit indicates the Upstream Fixed-Side Subsystem should reactivate main power and propagate the wake message Upstream. This may be accompanied by CMI_SSTART to indicate a staggered start behavior.

### 4.1.8. CMI_CLP_PRESENT

This bit is used to indicate the presence of the local Box. This bit is always set when sideband messages are sent. Other behaviors are beyond the scope of this specification.

### 4.1.9. CMI_HOT-PLUG_ATTENTION_BUTTON

This bit indicates that the CLP has detected a button press. Refer to Section 6 of the *PCI Express Base Specification* for more information

### 4.1.10. CMI_HOT-PLUG_INDICATOR_REQ

This bit indicates that the CLP has received direction to blink the indicator.

### 4.1.11. Cable Port Status Register

The Cable Port Status Register reflects the sideband message the BOX sent in the last sideband message. This allows to the CLP to verify the sideband status at any time.

It is strongly recommended that the CPSR be updated simultaneously when the transmission of the sideband message. If the sideband message is delayed, such as the BOX losing bus arbitration, the CPSR is still be updated.

### 4.1.12. Cable Port Capabilities Register

Each Fixed-side Subsystem contains a register space containing configuration information that is used to configure the system (see Table 4-3). These only need to be read by the Opposite Fixed-side, if Sideband Messages are supported, when the devices receive MGTPWR.

**Table 4-3. Fixed-side Capability Registers**

| Status Bit | Location | Description | Notes |
|---|---|---|---|
| 5 GT/s Support | 0 | Support for 5.0 GT/s operation | 5 GT/s operation is optional. Cable circuits may not support this data rate. |
| 8 GT/s Support | 1 | Support for 8.0 GT/s operation | This should be set to 1 |
| Reserved | 2 | | |
| Reserved | 3 | | |
| Reserved | 4 | | |
| SRIS Support | 5 | Support for Separate Reference Clocks with Independent Spread Spectrum | This is an optional feature. Both side must indicate support for this in the PCIe Cable Interface Capabilities Register. |
| Reserved | 6 to 15 | | |

# 4.2. Power Sequencing

Power sequencing for cabled PCI Express applications must comprehend multiple Boxes with independent power supplies, resets, and firmware. In addition, provisions are made to work with host systems that do not enable true power sequencing of the entire cabled system.

## 4.2.1. Power-Up Sequencing

By default, CPSR_RESET should be asserted by the management device when MGTPWR is applied. Downstream Subsystems should be powered before the Upstream Subsystem to allow the Downstream Subsystem time to fully initialize before the PCI Express link begins the discovery process.

The power-up sequence for the cable link and PCIe devices must follow the following sequence:

1. Enable MGTPWR, there is no precondition for enabling MGTPWR.
2. Initialize CMI Controller.
3. Set CPSR_CLP_PRESENT in the Fixed-Side Status register.
   Set CPSR_RESET for DSF Port(s).
   Local platform reset of PCI-Express devices must reflect the state of the sideband message CMI_RESET received on the USF Port if sideband messages are supported, otherwise it is governed by *PCI-Express Base Specification*.
4. Wait to proceed until assertion of CBLPRSNT#.
5. Read Cable Memory
   After delay of $T_{MGT\_INIT}$ of assertion of CBLPRSNT# as detected by the CMI Controller
   - Upstream Subsystem must attempt to read configuration data from A2h. If there is no response, the Upstream Subsystem must attempt to read from A0h.
   - Downstream Subsystems must read from A0h.

   See section 2.6 for more details regarding the configuration information. The CMI Controller is permitted to perform checks the integrity of the data before proceeding to the next step.

> **Observation:** The reading of the cable memory is necessary for proper configuration of the PCI Express device that interfaces to the cable port. However, there may be one or more failed attempts to read the data due to bus contention or non-response. The read process must be attempted until the configuration data is successfully acquired.

6. Read CLP Fixed-Side Memory if sideband messages are supported (based on the status of bit7 in offset 111 of the cable assembly's memory).

> **Note:** Bytes 2-18 must be read for configuration.

7. Prepare to Configure Core Clocking components and PCIe devices.
   a) The default behavior of Fixed-Side Subsystems is to:
      – Support SRNS. If the Fixed-Side cable port supports SRIS and the PCIe Cable Port Capabilities Register Bit 5 of the CLP indicates support of SRIS then the CMI controller must change the local reference clock to SRIS.
      – All physical cable port link widths and locations are set to a default configuration that is implementation dependent.
      – Cable link widths may link train to a lower link width, via normal PCIe Link Training protocols, if the cable assembly link width is smaller than the cable port link width. Lane 0 for any link must be connected to an effective Lane 0 of the CLP. Support of cable aggregation is not required.
   b) Port Width Configuration
      – Bifurcation of the Upstream Port
        CMI Controllers for DSF Ports that support cable bifurcation must obtain information that allows for port configuration from the CLP and the cable assembly before the PCI-Express link initialization begins (refer to *PCI Express Base Specification* for more details, Section 1).
        > DSF logical ports are permitted to be mapped to cabled ports in which the following fields of the CLP for each physical cable port match: Vendor ID, Device ID, Serial Number and the Connector Configuration field.
        > Bifurcation of a port, beyond an implementation's default configuration, with a cable that does not support sideband messages is beyond the scope of this specification.
        > Bifurcation of the USF Port is beyond the scope of this specification.
      – Cable Configuration
        Cabled links are permitted to aggregate multiple cable assemblies to form a single link. CMI Controllers for Ports supporting cable aggregation must be able obtain information that allows for port configuration from the CLP and the cable assembly before the PCI-Express hardware initialization begins. Cable aggregation with cable assemblies that do not support sideband messages is beyond the scope of this specification. Aggregation of Active Cable Assemblies is beyond the scope of this specification. The CMI controller compares the cable information across all physical ports connected to a logical port to determine if aggregation of the inserted cable assemblies is permitted

        If the cable assembly supports the maximum link, Cable memory byte 112 bits 0:2, of the logical port, further checks are not required.

        If any cable assembly inserted to the logical port is smaller than the logical port width, beginning with lowest lanes of the logical port, cables must contain matching data in the following bytes of the cable memory map unless otherwise specified:
        - Propagation Delay field
        - Common Port ID field
        - Cable Technology field
        - Cable Insertion Loss Attenuation fields

If the cables inserted to consecutive sets of x4 lanes of a logical cable port, identified by Port Order field of the CLP, do not meet matching criteria, then only the lanes connected to consecutive cable assemblies that meet the matching criteria, starting with the lowest lanes, are permitted to be enabled.

- It is recommended the CMI Controller use bits 0:1 of the Connector Configuration Register to verify aggregated cables are connected in the correct order such that the lanes of the Link are correctly ordered from 0-n to the other side.
- It is recommended that the PCI-Express device receivers of the disabled lanes be ignored to prevent unknown states or configurations within the PCI Express Device.
- Determine Equalizer presets and hints based on section 4.2.1.1. It is permitted for the presets and hints to be programmed at this stage if the implementation allows.

It is recommended to also change any settings that may be impacted by cable delay for both PCI-Express devices and the root complex. Examples of these settings may include, but are not limited to, replay timers, flow control credits, and presence detect circuit timings.

8. Enable Power Main.

Main power for any BOX with a Fixed-Side Subsystem may be enabled either via a sideband message or local power control. It is recommended that Boxes with local power control support initiation of staggered power sequencing over the cabled links to ensure all Downstream devices are ready for access when the Box with the root complex de-asserts Reset and to minimize potential strain on external power systems during the power-up process. The specific mechanism to enable this feature when using local power control is beyond the scope of this specification.

The CMI_SSTART bit in the sideband message enables the staggered start feature. This enables sequential power enablement of boxes in the cabled hierarchy. All Boxes with a cabled port must support the propagation of CMI_SSTART via sideband messaging. All boxes with an USF port must support the staggered start feature when directed via sideband message. Implementing staggered start via local control is strongly recommended.

There are two staggered start implementation options: DSF Gated and Local First. The DSF Gated implementation must send sideband messages with CMI_RESET Set, CMI_POWER_ENABLE Set, and CMI_SSTART Set to all DSF Ports, and then receive a sideband message with CMI_CLP_READY Set from all those DSF Ports before main power is permitted to be enabled. The Local First implementation permits local power to be enabled before sending sideband messages with CMI_RESET Set, CMI_POWER_ENABLE Set, and CMI_SSTART Set to all DSF Ports. With either method, all DSF ports must receive messages with CMI_CLP_READY Set on all of the USF Ports, or before de-asserting PERST# if the local Box contains the Root Complex.

**Implementation Note:** To enable further power staggering and reduce initial current draw, Boxes with multiple DSF ports may stagger sending the remote power message across its DSF ports that are not ready.

Sideband messages are not required to be sent to ports that show CBLPRESENT# asserted but the cable assembly does not support sideband messages.

> **Observation:** If the root complex does not receive a Successful Completion status to a request within 1.0 s of the Root Complex exiting Conventional Reset, the Root Complex is permitted to determine the Downstream BOX's PCIe device is broken. Usage of sideband messages enables power sequencing of the Boxes such that all Downstream Box's cabled ports are ready for configuration before the Root Complex begins sending Configuration Requests to cabled ports. Usage of cables that do not support sideband messages do not enable specific sequencing, Boxes power up in an undetermined sequence, and some or all Downstream Boxes may not be ready to meet time requirement for a Successful Completion.

Main power is permitted to be enabled for:

- Boxes with USF Port only
  Enable main power when the Box is directed to do so by either local control or if the USF port receives a sideband message with CMI_POWER_ENABLE Set.
- Boxes with DSF Port(s) only
  - If the Box receives a message with CMI_WAKE Set. If CMI_SSTART is also Set, the CMI controller implements the Gated First implementation of staggered start. The Local First implementation is permitted if the CMI controller prevents platform reset de-assertion until the Downstream Boxes are ready.
  - If directed by local power control to enable main power,
    - If either the cable inserted to the DSF port does not support sideband messages or the staggered start feature is not to be initiated, then main power is permitted to be enabled immediately.
    - If the staggered start feature is to be initiated, the CMI controllers implements the Gated First implementation of staggered start. The Local First implementation is permitted if the CMI controller prevents platform reset de-assertion until the Downstream Boxes are ready.

> **Implementation Note:** Cabled ports provided via an adapter card, or module, may not have the ability to gate power for the Box in which the adapter is inserted. A power switch may be provided by the adapter to have the CMI controller initiate the remote power enable feature and once the DSF port(s) receive messages with CMI_CLP_READY, the local wake signal for the box may be asserted according to form factor requirements.

> **Observation:** The Box with only a DSF Port contains the root complex. If sideband messages are not supported by the cable, then Downstream Boxes may not configure properly before the Box de-asserts platform reset, unless the Downstream Boxes are power on first by other means.

- BOX's with both a USF Port and at least one DSF Port.

**Implementation Note:** Either Gated DSF or Local First are permitted to be implemented by this type of box.

  - If directed by local power control
    - If either the cable inserted to the DSF port does not support sideband messages or the staggered start feature is not to be initiated, then main power is permitted to be enabled immediately.
    - If staggered start is to be initiated, either method is permitted.
  - Directed via a DSF Port by receiving a sideband message with CMI_WAKE Set
    - If the sideband message also has CMI_SSTART Set, either staggered start implementation is permitted.
    - If the sideband message also has CMI_SSTART Cleared, main power is permitted to be enabled, and sideband messages must be sent to all other DSF Ports in the box with both CMI_POWER_ENABLE and CMI_RESET Set, but CMI_SSTART Cleared. The sequencing of sending this message to the DSF ports is not specified. A message with CMI_WAKE Set is permitted to be sent to the USF Port.
  - Directed via USF Port be receiving a sideband message with CMI_POWER_ENABLE Set
    - If the received sideband message also includes CMI_SSTART Set, either staggered start implementation is permitted.
    - If the received sideband message includes CMI_SSTART Cleared, a sideband message with both CMI_RESET and CMI_POWER_ENABLE Set, and with CMI_SSTART Cleared, must be sent to all DSF Ports with CBLPRSNT# asserted; and is permitted to enable main power immediately.

9. Set CPSR_CLP_READY when main power is within operating limits, clocks are stable, cable memory configuration data has been obtained, and the box is ready for CMI_RESET to be de-asserted.

**Implementation Note:** Implementations of a Box are permitted to have additional criteria gate CPSR_CLP_READY. These may include initialization of functions outside the PCI Express specific functions.

10. Indicate to Cable Link Partner that box is ready.
    - If power was initiated either by receiving a message with CMI_WAKE Set on a DSF Port set, or by local power control
      - Send a message to the USF port with both CMI_WAKE and CMI_CLP_READY Set. CMI_SSTART must reflect the state of CMI_SSTART from the message on the DSF port with CMI_WAKE Set, or the staggered start implementation for the local power control mechanism.
      - Send a message to all DSF ports with CMI_CLP_READY Set.
    - If power was initiated by receiving a message with CMI_POWER_ENABLE Set from the USF Port
      - Send a message to the USF port with CMI_CLP_READY. CMI_SSTART must reflect the state of CMI_SSTART from the message on the USF port with CMI_POWER_ENABLE Set.

11. CMI_WAKE is to be Cleared from CMI messages once a message is received with CMI_RESET Cleared.

#### 4.2.1.1. Equalizer Settings Based on Cable Memory

Cabled PCIe ports need to be configured to work correctly in the cabled environment. Loss for the cable is stored in the cable memory map. The Fixed-side CMI controller sets the transmitter and receiver settings for 8GT/s based on loss value for 4 GHz. The recommended settings for cable port transmitter and receiver starting values are listed in Table 4-4. These values are derived from passive cable assemblies.

**Table 4-4. Cable Port Transmitter and Receiver Starting Values**

| Cable Signal Transmission (in dB) | Receiver CTLE Settings |
|---|---|
| -3 or higher [1, 2] | -6 |
| -4 | -7 |
| -5 | -7 |
| -6 | -8 |
| -7 | -8 |
| -8 | -9 |
| -10 | -10 |
| -11 | -11 |
| -12 or lower [3] | -12 |

**Notes:**

1. Cable assemblies reporting 0 dB are electrically buffering the signal before exiting the cable assembly and indicate a full-swing signal is being sent from the cable assembly to the Fixed-side Subsystem.
2. It is recommended that inserted cable assemblies reporting a higher insertion loss value that -3 dB be set to the minimum receiver value of -6 dB. Receivers that support Rx CTLE values of less than -6dB are beyond the scope of this specification.
3. It is recommended that inserted cable assemblies reporting a lower insertion loss value that -12 dB be set to the maximum receiver value of -12 dB. Receivers that support Rx CTLE values of less than -12 dB are beyond the scope of this specification.

It is recommended that transmitters begin with a setting of P6 unless the loss of the cable is reported as 0dB. A setting of 0dB implies a fully buffered cable and in this case and it is recommended the preset be configured for P4.

Cable technology also impacts the presets. Byte 147 indicates the cable technology. Values of Ah, Bh, or Fh indicates the total effective loss of the cable. Values of 0h or Ch indicates that the cable is fully buffered. It is recommended that ports connected to fully buffered cable assemblies set the Tx equalization to P3 and Rx to its minimum value. A value of Dh indicates the Rx end is fully buffered (or non-linear), but the Tx end is not.

The receiver should be set to the minimum equalization setting while the transmitter should optimize for the loss of the transmitter PCB interconnect and the reported cable loss. Refer to the *PCIe Base Specification* for the definitions of the presets.

The method for the CMI controller to program the PCIe component is implementation specific and beyond the scope of this specification.

#### 4.2.1.2. PCI Express Configurations Based on Cable Memory

Long cable delays impact the link performance and the PCI Express protocol. Replay Timers and Completion Timers are tripped if acknowledgements are not received in an expected timeframe. Limited credit allocation may limit the amount of data in flight over long cable reaches. The Propagation Delay field provides information that is used to set timers and credits for optimal performance. It is recommended that if the cable technology is reported as 0h, then the Completion Timeout field in the Device Control 2 Register be set to Range B or higher. Devices should be programmed with the maximum number of credits possible to allow for as many packets to be in flight as possible. For any type of active cable, indicated by the Cable Technology field, it is recommended to set N_FTS to 255 and the Extended Sync bit to 1, in the Link Control Register, to allow maximum time for link training. It is also recommended that ASPM be disabled for all cable technologies except passive cables to avoid gain and noise problems on the link.

Need Image

**Figure 4-1. Power-up Sequencing**

### 4.2.2. Power-Down Sequencing

Figure 4-2 illustrates the power-down sequencing resulting from a graceful or Hot-Unplug event.

#### 4.2.2.1. Sequence Steps

Following are the sequence steps:

- PCI-Express lanes transitions to inactive state (Device in D3hot)

- A sideband message is initiated by the BOX with the Root Complex with CMI_RESET Set. The message is propagated from all USF Ports to any DSF Ports immediately, taking priority over any other sideband messages that may be queued.
- CMI controller for a BOX with USF Port(s) asserts local reset signal.
- A sideband message is initiated by the BOX with the Root Complex with both CMI_POWER_ENABLE and CMI_CLP_READY Cleared. The message is propagated from all USF Ports to any DSF Ports.
- CMI controller for a BOX with USF Port(s) is permitted to disable main power

> **Observation:** Most Boxes are expected to follow the status of CMI_POWER_ENABLE that originates in the BOX with the root complex, after initial Reset de-assertion. There are implementations in which it is desirable to leave the Box powered without PCI-Express communication. Boxes with USF ports are permitted to leave main power enabled when a sideband message is received with CMI_POWER_ENABLE Cleared. However, the PCI-Express device must follow the state of CMI_RESET.

#### 4.2.2.2. Timings TBD

Need Image

**Figure 4-2. Power-Down Sequencing**

> **Note:** It is recommended to power off the Downstream Subsystem after the Upstream Subsystem, if using local power control, to allow system software to gracefully exit before the hardware is effectively removed from the system topology.

### 4.2.3. Cable Management Interface Timings

The operational description of the CMI are specified by the *PCI Express OCuLink Specification*, Appendix B. General protocol and Read/Write operations are specified. In addition to these parameters, the arbitration protocol from the I2C Revision 2.1 is followed to support the multi-master nature of the CMI and potential bus contention between the Fixed-Side CMI controllers at each end of a cabled link.

### 4.2.4. Link Power Management

The *PCI Express Base Specification* specifies Link Power Management states, some of which are required by all existing form factor specifications while others may be optional. A cabled PCI Express Port, adhering to this revision of the specification, may be implemented on top of any other form factor. Support of Link Power Management states are considered to be a system level requirement and are not imposed by this specification.

Downstream Subsystems should support all defined Link Power Management states to guarantee compliancy with system level specifications.

Support for dynamic control of reference clock via Clock Power Management during L1 states, as described in the *PCI Express Base Specification*, is not provided by this specification.

**Note:** Some active cables do not respond well to lack of signal transitions during the low power states defined by the PCI Express Base Specification. Systems should disable the Link Power Managing states to increase compatibility and improve performance.

### 4.2.5. Miscellaneous

DCP, SLOT PWR, SLOT CLOCK, Hierarchy ECN, Aux Field, DRS/DFS, Power Budgeting.

1108

5

# 1109 5. Fixed-Side Connector Definition

1110 The SFF-8644 specification defines distinct connectors to support multiples of x4 Link width. Links
1111 with lower lane counts must use this connector and have two or more unused Lanes. Links greater
1112 than four different Link widths: x1, x4, x8, x12, and x16 Lanes use multiple connectors to achieve
1113 the total Link width. One cable connector is defined for each of these four Link widths. A low cost
1114 x2 connector may be defined in a future specification revision. All electrical and mechanical testing
1115 is to be conducted in compliance with EIA-364.

## 1116 5.1. Signal Description

1117 The external PCI Express cable connector and cabling support the following signals:

1118 ❑ PETpN/PETnN (required): PCI Express Transmitter pair(s), labeled where N is the Lane
1119 number (starting with 0); "p" is the true signal while "n" is the complement signal.

1120 ❑ PERpN/PERnN (required): PCI Express Receiver pair(s), labeled where N is the Lane number
1121 (starting with 0); "p" is the true signal while "n" is the complement signal.

1122 ❑ Auxiliary signals as defined in Chapter 2, Auxiliary Signals.

1123 The Fixed-side PCI Express signals must be connected to the PCI Express device, transmitters to
1124 the PETxY pins and receivers to the PERxY pins as described in Table 6-2. The cable assembly
1125 provides the cross over path that enables transmit pins to send data to receive pins.

# 5.2. Fixed-Side Connector

Refer to the *SFF-8644 Specification* for mechanical information.

## 5.2.1. Pin-out

Table 5-1 contains the pinout of the base connector. If the implementation utilizes more than a single x4 fixed connector, the pinout should follow that shown in Table 5-2. Figure 5-2 shows the four Fixed-side x4 connector pin assignments to enable a x16 link. Figure 5-3 illustrates the ordering of the connectors on a card edge.

The connector pin definitions are referenced to the PCI Express device pins of the same Fixed-side Subsystem. PETxY pin of the connector must be connected to the transmitter pins of the device. PERxY pins of the connector must be connected to the receiver pins of the PCI Express device.

**Note:** The PCI Express External Cabling SFF-8644 fixed-side connector pinout lane mapping, although very similar, is not the same as the SAS SFF-8449 pinout lane mapping. Although the cables are compatible, attention is drawn to the fact that lanes 0 and 1 are swapped in the PCI Express definition. The SFF-8644 lane configuration facilitates simplified routing lane ordering which is not important for formation of SAS wide ports. It is however, critical to PCI Express in order to enable lanes to be combined into x4, x8, and x16 links.

**Observation:** The connector used for PCI-Express is also used by other common technologies; however, the pinouts are not the same. Designers should use caution when wiring the PCI-Express lanes 0 and 1 to avoid inverting the lane sequence.

### Table 5-1. Fixed-side Connector Pinout

| | Column | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **Row** | **9** | **8** | **7** | **6** | **5** | **4** | **3** | **2** | **1** |
| **D** | GND | PETn2 | PETp2 | GND | PETn1 | PETp1 | GND | MGTPWR | PWR |
| **C** | GND | PETn3 | PETp3 | GND | PETn0 | PETp0 | GND | CMISDA | CMISCL |
| | | | | | | | | | |
| **B** | GND | PERn2 | PERp2 | GND | PERn1 | PERp1 | GND | CBLPRSNT# | PWR |
| **A** | GND | PERn3 | PERp3 | GND | PERn0 | PERp0 | GND | CINT# | CADDR |

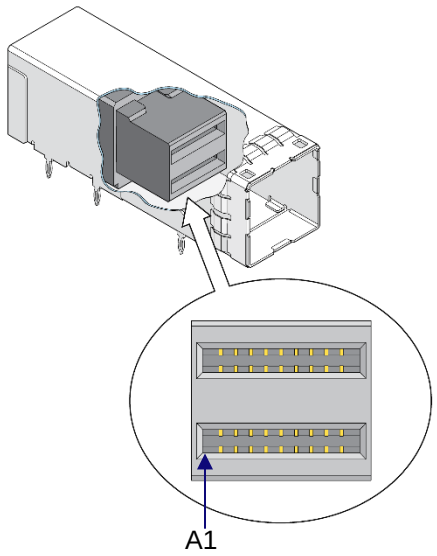


1143
1144

1145 **Figure 5-1. SFF-8644 Fixed-side Connector in Fixed-side Cage**

1146 **Table 5-2. Fixed-side Multi-connector Lane**

| | | Column | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | Row | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| Connector 0 | D | GND | PETn2 | PETp2 | GND | PETn1 | PETp1 | GND | MGTPWR | PWR |
| | C | GND | PETn3 | PETp3 | GND | PETn0 | PETp0 | GND | CMISDA_0 | CMISCL_0 |
| | B | GND | PERn2 | PERp2 | GND | PERn1 | PERp1 | GND | CBLPRSNT#_0 | PWR |
| | A | GND | PERn3 | PERp3 | GND | PERn0 | PERp0 | GND | CINT#_0 | CADDR_0 |
| Connector 1 | D | GND | PETn6 | PETp6 | GND | PETn5 | PETp5 | GND | MGTPWR | PWR |
| | C | GND | PETn7 | PETp7 | GND | PETn4 | PETp4 | GND | CMISDA_1 | CMISCL_1 |
| | B | GND | PERn6 | PERp6 | GND | PERn5 | PERp5 | GND | CBLPRSNT#_1 | PWR |
| | A | GND | PERn7 | PERp7 | GND | PERn4 | PERp4 | GND | CINT#_1 | CADDR_1 |
| Connector 2 | D | GND | PETn10 | PETp10 | GND | PETn9 | PETp9 | GND | MGTPWR | PWR |
| | C | GND | PETn11 | PETp11 | GND | PETn8 | PETp8 | GND | CMISDA_2 | CMISCL_2 |
| | B | GND | PERn10 | PERp10 | GND | PERn9 | PERp9 | GND | CBLPRSNT#_2 | PWR |
| | A | GND | PERn11 | PERp11 | GND | PERn8 | PERp8 | GND | CINT#_2 | CADDR_2 |
| Connector 3 | D | GND | PETn14 | PETp14 | GND | PETn13 | PETp13 | GND | MGTPWR | PWR |
| | C | GND | PETn15 | PETp15 | GND | PETn12 | PETp12 | GND | CMISDA_3 | CMISCL_3 |
| | B | GND | PERn14 | PERp14 | GND | PERn13 | PERp13 | GND | CBLPRSNT#_3 | PWR |
| | A | GND | PERn15 | PERp15 | GND | PERn12 | PERp12 | GND | CINT#_3 | CADDR_3 |

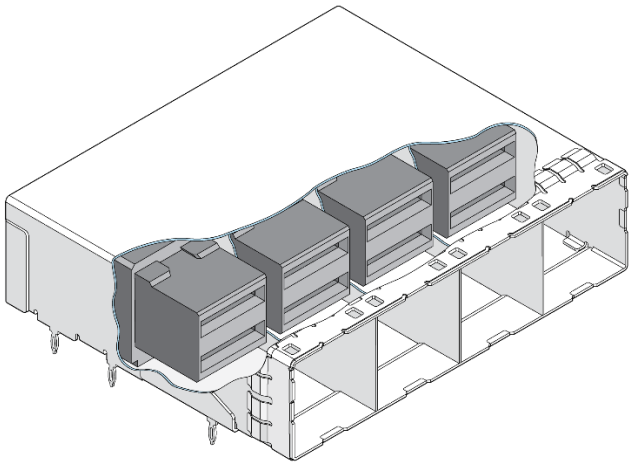

**Figure 5-2. Four Fixed-side x4 Connector Configuration**

Connectors must expand to the right of connector 0 when viewed from the external connector face, component side up, as shown in Figure 5-3.

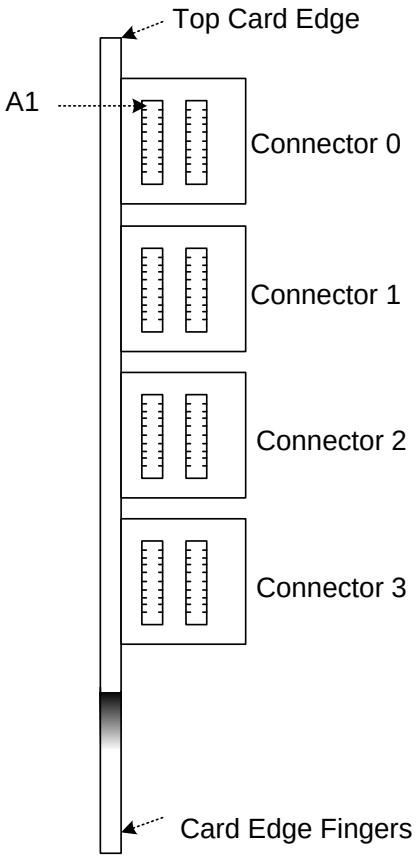


**Figure 5-3. Temp Image, Connector Orientation Diagram**

6

# 6. Cable Specification

## 6.1. SFF-8644 Cable Assembly

There are two classes of cable assemblies defined for this specification.

- PCIe Sideband-Isolated
  The first class of cable has an isolated management interface with regard to each end of the cable assembly. In this class, there is no mechanism for external sidebands. Examples of this cable are assemblies that conform to the *SFF-8449 Specification* used in storage applications, or fiber optic cables that do not provide a mechanism to effectively send CMISDA and CMISCL across the assembly. The SFF-8449 pin definitions are different than PCI Express definitions, though the cables are functionally compatible. See section 5.2.1 for further details.
- PCIe Sideband-Enabled
  The second class of cable, or PCIe Gen3 cable, links the management interfaces at either end of the assembly together. This enables either end of the link to read the memory in the cable and for each end to communicate with the other. With the fixed ends linked via CMI, full sideband support is enabled with a higher level of configuration and other features supported.

The cable assembly must follow the pinout defined in Table 6-1. The wiring diagram in Table 6-2 indicates the internal cable connectivity from end to end. Where the connectivity does not show a direct arrow, the letters indicate which end the pin connects to at the opposite end. Power is not provided to the opposite end of the cable assembly.

Active cables must include DC-blocking capacitors between the active component and the receive pin of the connector of the cable assembly. Active cable assemblies must preserve the in-band presence detect mechanism of the SERDES as specified in the PCI Express Base Specification.

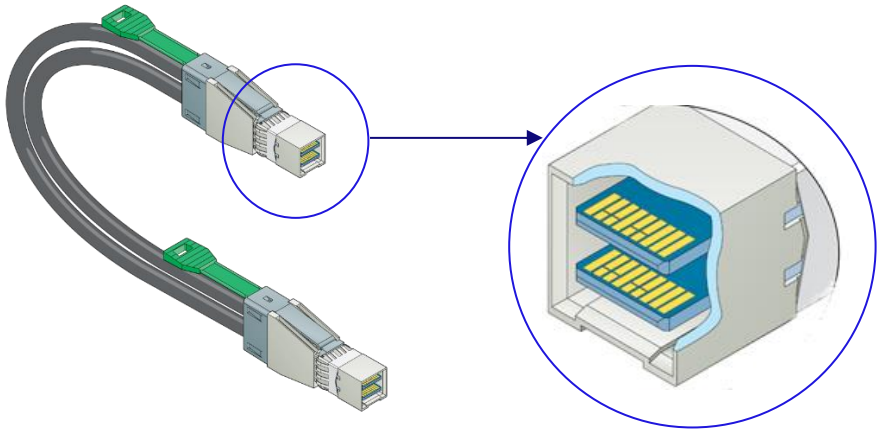

**Figure 6-1. Cable Assembly with SFF-8644 Connectors**

**Table 6-1. Connector Pin assignments**

| Row | Column 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
|---|---|---|---|---|---|---|---|---|---|
| **D** | GND | PETn2 | PETp2 | GND | PETn1 | PETp1 | GND | MGTPWR | PWR |
| **C** | GND | PETn3 | PETp3 | GND | PETn0 | PETp0 | GND | CMISDA | CMISCL |
| | | | | | | | | | |
| **B** | GND | PERn2 | PERp2 | GND | PERn1 | PERp1 | GND | CBLPRSNT# | PWR |
| **A** | GND | PERn3 | PERp3 | GND | PERn0 | PERp0 | GND | CINT# | CADDR |

**Table 6-2. Cable Wiring**

| Side 1 | | | Side 2 | |
|---|---|---|---|---|
| CADDR | A1 | ← PC[1] x[2] → | C1 | CMISCL |
| CINT# | A2 | ← PC[1] y[2] → | C2 | CMISDA |
| GND | A3 | ← → | C3 | GND |
| PERp0 | A4 | ← → | C4 | PETp0 |
| PERn0 | A5 | ← → | C5 | PETn0 |
| GND | A6 | ← → | C6 | GND |
| PERp3 | A7 | ← → | C7 | PETp3 |
| PERn3 | A8 | ← → | C8 | PETn3 |
| GND | A9 | ← → | C9 | GND |
| PWR | B1 | No Wire | D1 | PWR |
| CBLPRSNT# | B2 | No Wire | D2 | MGTPWR |
| GND | B3 | ← → | D3 | GND |
| PERp1 | B4 | ← → | D4 | PETp1 |

| Side 1 | | | Side 2 | |
|---|---|---|---|---|
| PERn1 | B5 | ← → | D5 | PETn1 |
| GND | B6 | ← → | D6 | GND |
| PERp2 | B7 | ← → | D7 | PETp2 |
| PERn2 | B8 | ← → | D8 | PETn2 |
| GND | B9 | ← → | D9 | GND |
| | | | | |
| CMISCL | C1 | ← x[2] PC[1] → | A1 | CADDR |
| CMISDA | C2 | ← y[2] PC[1] → | A2 | CINT# |
| GND | C3 | ← → | A3 | GND |
| PETp0 | C4 | ← → | A4 | PERp0 |
| PETn0 | C5 | ← → | A5 | PERn0 |
| GND | C6 | ← → | A6 | GND |
| PETp3 | C7 | ← → | A7 | PERp3 |
| PETn3 | C8 | ← → | A8 | PERn3 |
| GND | C9 | ← → | A9 | GND |
| PWR | D1 | No Wire | B1 | PWR |
| MGTPWR | D2 | No Wire | B2 | CBLPRSNT# |
| GND | D3 | ← → | B3 | GND |
| PETp1 | D4 | ← → | B4 | PERp1 |
| PETn1 | D5 | ← → | B5 | PERn1 |
| GND | D6 | ← → | B6 | GND |
| PETp2 | D7 | ← → | B7 | PERp2 |
| PETn2 | D8 | ← → | B8 | PERn2 |
| GND | D9 | ← → | B9 | GND |

**Notes:**

1. PC signals are only connected to the near-end backshell and paddle card. They have no wire traversing the cable assembly.
2. **x** and **y** indicate that the signals connect by wire to another location relative to the bulk wire for side-band enable cables. For cables that are not side-band enabled, these signals do not have cable wires and only connect to the near-end paddle card.

1182

**Note:** There is no specific x1 cable assembly for this specification. Links smaller than x4 must use the x4 cable assembly and the link must downtrain. For ports smaller than x4, the port must begin with Lane 0.

1183

1184 Refer to the *SFF-8644 Specification* for the mechanical requirements. Electrical performance is
1185 discussed in section 6.2. Cable assemblies must provide the information specified in the memory
1186 map described in Table 6-3 and Table 6-4.

# 6.2. Cable Management Interface

Cable assemblies operating at 8.0 GT/s influence the link equalization process since a predetermined equalization setting is not compatible with all cable assemblies. The cable assembly must include a memory location to provide, at a minimum, the loss information for the cable such that both the Upstream and Downstream Subsystems may derive appropriate equalization starting points for setting the Transmitter and Receiver Presets. The specifics for the memory map are located in 6.2.2.3.

The management interfaces are limited to four electrical loads and four addresses. These include a device in the Upstream Fixed-side portion of the link, the Downstream Fixed-side portion of the link, and each end of the cable assembly. These must be isolated from any other management interfaces in the system to limit conflicts and incompatibilities.

Monolithic cable assemblies (those designed to be larger than x4) need only provide one set of sideband signals, except power, to the Subsystem. The sidebands must be in the connector in which Lane 0 is housed. Manufacturers need to ensure that the loss information is consistent across the entire assembly. For cable assemblies that provide for monitoring of the cable components for cable management and monitoring connected devices, these cable assemblies need to implement the Cable Management Interface lines for each x4 connector.

For cables that exceed the capacitive loading limits given in Table 2-2, the cable assembly may need to provide mechanisms to ensure the rise time requirements are met.

There are circuits and components required to support functions of the cable. These are defined per sideband, see Chapter 2.

## 6.2.1. Memory Map

## 6.2.2. Memory Map Architecture

This section defines the non-volatile memory map and protocol utilized for external cable identification and management. The cable memory address is specified in section 2.4, CADDR. The LSB of the memory is configurable based on the method described. However, cable assemblies supporting side-band messages need only have configuration data stored at address A0h and this does not need to be configurable via CADDR.

Unless otherwise stated, all informative fields must contain accurate data. Any used to represent a field that is unspecified is invalid. All reserved bytes and bits must be filled with logical zeroes.

For future expansion, it is recommended to keep a provision for the upper bytes 128-255 to be swapped in multiple pages, as is conventional in other memory map implementations. Which actual memory bytes are seen in addresses 128-255 then depends on the page number written in byte 127. Figure 6-2 shows the structure of the PCI Express cable CMI memory map.

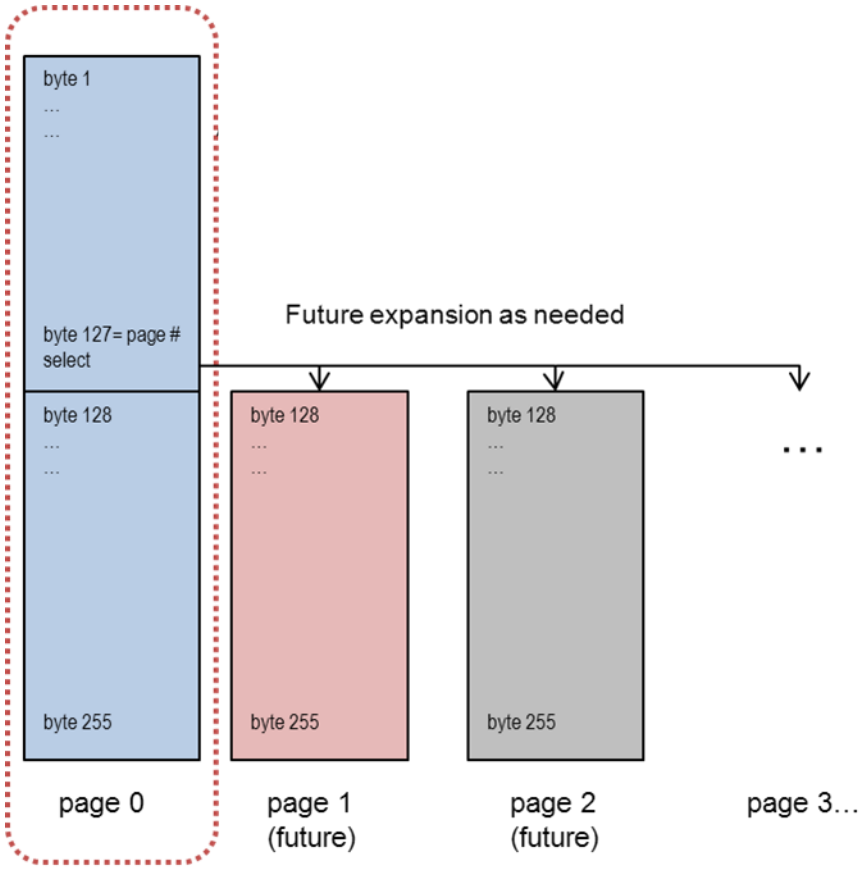


1221

1222 **Figure 6-2. Structure of the PCI Express Cable CMI Memory Map**

1223 The most significant bit (MSB) in a byte is bit 7, and the least significant bit (LSB) is bit 0. For fields
1224 that are two bytes long, the MSB is bit 7 of the higher byte number, and the LSB is bit 0 of the
1225 lower byte number.

1226 Fields are defined as Read Only (RO), Write Only (WO), or Read/Write (R/W).

### 1227 6.2.2.1. Bytes 0-127 (Lower Memory)

1228 Table 6-3 lists lower memory fields for Page 0. Fields are defined as Read Only (RO), Write Only
1229 (WO), or Read/Write (R/W).

## Table 6-3. Page 0 Lower Memory Fields

| Byte | Description | Values | R/W | Notes |
|---|---|---|---|---|
| 0 | Cable Identifier | 0Fh – x4,<br>10h – x8,<br>00h – x16 | RO | SFF-8636 cables use 00h as an undefined cable type. PCIe uses this in combination with Byte112, bit 0:2 to indicate a x16 PCIe cable. |
| 1 | Reserved | | | |
| 2 | Status Indicators | Bits 0 - Data_Not_Ready<br>Bit 1 – Int#<br>Bit 2 - Flat_mem<br>Bits - 3:7 | RO | See 6.2.2.3.1 |
| 3 | Interrupt Flags- LOS | For all bits, a value of 0 indicates deasserted.<br>Bit 0 – Rx0<br>Bit 1 – Rx1<br>Bit 2 – Rx2<br>Bit 3 – Rx3<br>Bit 4 – Tx0<br>Bit 5 – Tx1<br>Bit 6 – Tx2<br>Bit 7 – Tx3 | RO | 6.2.2.3.2, Optional |
| 4 | Reserved | | | |
| 5 | Reserved | | | |
| 6 | Interrupt Flags – Temp Alarm | Bits 0:3 – Reserved<br>Bit – 4 Temp Low Warning<br>Bit – 5 Temp High Warning<br>Bit – 6 Temp Low Alarm<br>Bit – 7 Temp High Alarm | RO | 6.2.2.3.3, Optional |
| 7 | Interrupt Flags – Voltage Alarm | Bits 0:3 – Reserved<br>Bit – 4 VCC Low Warning<br>Bit – 5 VCC High Warning<br>Bit – 6 VCC Low Alarm<br>Bit – 7 VCC High Alarm | RO | 6.2.2.3.4, Optional |
| 8-21 | Reserved | | | |
| 22 | Module Monitors – Temperature MSB | | RO | Optional |
| 23 | Module Monitors – Temperature LSB | | RO | Optional |
| 24 | Reserved | | | |
| 25 | Reserved | | | |
| 26 | Module Monitors – Supply Voltage MSB | | RO | Optional |
| 27 | Module Monitors – Supply Voltage LSB | | RO | Optional |

| Byte | Description | Values | R/W | Notes |
|---|---|---|---|---|
| 28-107 | Reserved | | | |
| 108 | Propagation delay | Electrical Delay in ns. Upper 8 bits of 16-bit value. | RO | See 6.2.2.3.5 |
| 109 | Propagation delay | Electrical Delay in ns. Lower 8 bits of 16-bit value. | RO | See 6.2.2.3.5 |
| 110 | Reserved | | | |
| 111 | Cabled PCIe Capabilities 1 | Data Rate Identifier<br>Bit 0 – 2.5 GT/s Data Rate Supported. Must be set to 1b.<br>Bit 1 – 5.0 GT/s Data Rate Supported<br>Bit 2 – 8.0 GT/s Data Rate Supported.<br>Bits 3:4 – Reserved<br>Bit 5 - Legacy Port Adapter<br>Bit 6 – Reserved.<br>Bit 7 – Support for Sideband Messages | RO | See 6.2.2.3.6 |
| 112 | Cabled PCIe Capabilities 2 | Bits 0:2<br>000b – x1<br>001b – x2<br>010b – x4<br>011b – x8<br>100b – x12<br>101b – x16<br>All other values Reserved<br>Bits 3:7 – Reserved | RO | See 6.2.2.3.7 |
| 113-118 | Reserved | | | |
| 119-122 | Password Change | | WO | 6.2.2.3.8 |
| 123-126 | Password Entry | | WO | 6.2.2.3.8 |
| 127 | Page Select Byte | | R/W | 6.2.2.3.9 |

### 6.2.2.2. Bytes 128-255 (“Page 0, Upper Memory”)

Table 6-4 lists the Upper Memory Fields for Page 0.

**Table 6-4. Page 0 Upper Memory Fields**

| Byte | Description | Values | R/W | Notes |
|---|---|---|---|---|
| 128 | Cable Identifier | Must match byte 0 | RO | Must match Byte 0 |
| 129 | Reserved | | RO | |
| 130-146 | Reserved | | | |
| 147 | Cable Technology | Bits 0:3 Reserved<br>Bites 4:7 Cable Technology Type | RO | Cable equalization and medium<br>See 6.2.2.3.10 |
| 148-163 | Vendor Names | ASCII string (16 characters) | RO | 6.2.2.3.11 |
| 164 | Reserved | | | |
| 165-166 | PCISIG Vendor ID | Hex Number (two bytes) | RO | 6.2.2.3.12 |
| 167 | Reserved | | | |
| 168-183 | Vendor Part Number | ASCII string (16 characters) | RO | 6.2.2.3.13 |
| 184-185 | Vendor Revision | ASCII string (2 characters) | RO | 6.2.2.3.14 |
| 186 | Attenuation 1.25 GHz | Hex number | RO | 6.2.2.3.15 |
| 187 | Attenuation 2.5 GHz | Hex number | RO | 6.2.2.3.15 |
| 188 | Attenuation 4.0 GHz | Hex number | RO | 6.2.2.3.15 |
| 189 | Attenuation 8.0 GHz | Hex number | RO | 6.2.2.3.15 |
| 190 | Maximum case temperature | Hex number | RO | 6.2.2.3.16 |
| 191 | Checksum - Base | Checksum on Bytes 128-190 | RO | 6.2.2.3.17 |
| 192 | Reserved | | | |
| 193 | Reserved | | | |
| 194 | Reserved | | | |
| 195 | Options | Bits 0:7 – Reserved | RO | 6.2.2.3.18 |
| 196-211 | Vendor Serial Number | ASCII string (16 characters) | RO | 6.2.2.3.19 |
| 212-217 | Vendor Data Code | ASCII string (yymmdd) | RO | 6.2.2.3.20 |
| 218-219 | Vendor Lot Code | | RO | 6.2.2.3.21 |
| 220-222 | Reserved | | | |
| 223 | Checksum - Extended | Checksum for Bytes 192-222 | RO | 6.2.2.3.22 |
| 224-255 | Vendor Specific Information | Vendor Specific (32 Bytes) | RO | 6.2.2.3.23 |

## 6.2.2.3. Field Descriptions

This section defines the non-volatile memory map and protocol utilized for external cable identification and management. All reserved bytes must be filled with logical zeroes.

All bits are asserted in the as a positive value.

### 6.2.2.3.1. Byte 2, Status Indicators

The Status Indication field indicates the status of cable management information. Table 6-5 lists the bit descriptions.

**Table 6-5. PCI Express Cable Status Indicator**

| Bit | Description |
|---|---|
| 0 | Data_Not_Ready: Asserted until a valid set of monitor readings |
| 1 | Int# (Optional): Digital indicator reflecting CINT# |
| 2 | Flat_mem |
| 3:7 | RsvdP |

### 6.2.2.3.2. Byte 3, Interrupt Flags – LOS

The Byte 3 latched bits indicate a loss of signal per lane.

The asserted value must be 1b. Once a bit is Set, CINT# is asserted and the value is retained until a read operation of the bit or the free side is reset by removal and assertion of MGTPWR. The bit is permitted to be set again any time the criteria for the assertion of the signal is met.

Bits 0 to 3 indicate the Rx lanes, with 0 being Lane 0 and the bit location incrementing per lane.

Bits 4 to 7 indicate the Tx lanes, with bit 4 being Lane 0 and the bit location incrementing per lane.

### 6.2.2.3.3. Byte 6, Interrupt Flags – Temp Alarm

The Byte 6 latched bits indicate a temperature condition.

See section 6.2.2.3.2 for the ordering of the bits.

### 6.2.2.3.4. Byte 7, Interrupt Flags – Voltage Alarm

The Byte 7 latched bits indicate a voltage condition.

See section 6.2.2.3.2 for the ordering of the bits.

### 6.2.2.3.5. Bytes 108-109, Propagation Delay

These 2 bytes indicate the range of the one-way propagation delays through the cable assembly of all the signal pairs. The propagation delays all fall between a lower bound and an upper bound. The lower bound is defined as the largest multiple of 250 ps that is less than or equal to the propagation delay of the fastest signal pair in the cable. In other words, it is the smallest propagation delay, rounded down to a multiple of 250 ps. The upper bound is defined as the smallest multiple of 250

ps that is greater than or equal to the propagation delay of the slowest signal pair in the cable. In other words, it is the largest propagation delay, rounded up to a multiple of 250 ps.

These upper and lower bounds are expressed by a combination of 2 parameters, Base and Span, in the Propagation delay field, as shown in Figure 6-3.

| Byte | 108 | | | | | | | 109 | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Bit | 7 (msb) | 6 | 5 (lsb) | 4 (msb) | 3 | 2 | 1 | 0 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 (lsb) |
| Contents | Span | | | Base | | | | | | | | | | | | |

## Figure 6-3. Bytes 108-109, Propagation Delay

These parameters are defined by the following equations:

- Base = (Lower Bound) / (250 ps)
- Span = 0, for (Upper Bound - Lower Bound) ≤ 250 ps
- Span = (Upper Bound - Lower Bound) / (250 ps) - 2, for (Upper Bound - Lower Bound) > 250 ps

The Lower Bound, Upper Bound, Base, and Span values for some example cables are shown in the following table:

## Table 6-6. Lower Bound, Upper Bound, Base, and Span Values

| Parameter | Cable 1 | Cable 2 | Cable 3 | Cable 4 |
|---|---|---|---|---|
| Fastest Pair's Propagation Delay | 10.2 ns | 133.6 ns | 794.8 ns | 2047.9 ns |
| Slowest Pair's Propagation Delay | 12.1 ns | 133.7 ns | 795.2 ns | 2049.2 ns |
| Lower Bound | 10.00 ns | 133.50 ns | 794.75 ns | 2047.75 ns |
| Upper Bound | 12.25 ns | 133.75 ns | 795.25 ns | 2049.25 ns |
| Upper Bound - Lower Bound | 2.25 ns | 0.25 ns | 0.5 ns | 1.5 ns |
| Base | 40d (28h) | 534d (214h) | 3179d (C6Bh) | 8191d (1FFFh) |
| Span | 7h (111b) | 0h (000b) | 0h (000b) | 4d (100b) |

A cable management controller computes the Upper and Lower Bound for a cable based on the following equations:

- Lower Bound = Base * 250 ps
- Upper Bound = (Base + Span + 2) * 250 ps

**Observation:** The lowest Upper Bound that is expressed by the Span parameter is 500 ps more than the Lower Bound.

**Observation**: Because actual propagation delays do not fall on exact 250 ps boundaries, virtually all cables report a range of propagation delays that is up to 250 ps longer than the actual cable skew. Thus, a cable that has almost 2 ns of skew most likely report a 2.25 ns range around its fastest and slowest propagation delays. Because of this, a cable management controller should not compare a cable's own reported Upper Bound and Lower Bound with each other when determining cable compatibility for port aggregation. It should compare each cable's reported Upper Bound with every other cable's reported Lower Bound. Any difference greater than +2 ns indicates that the combination has excessive skew and is not aggregated.

#### 6.2.2.3.6. Byte 111, Cabled PCIe Capabilities 1

Table 6-7 lists the bit descriptions for the PCI Express cable port capabilities, Register 1.

**Table 6-7. PCI Express Cable Port Capabilities Register 1**

| Bit | Description |
|---|---|
| 0 | 2.5 GT/s, must be set to 1b |
| 1 | 1b if 5.0 GT/s capable, else 0b |
| 2 | 1b if 8.0 GT/s capable, else 0b |
| 3-4 | RsvdP |
| 5 | Legacy Port Adapter |
| 6 | RsvdP |
| 7 | Support for Sideband Messages |

#### 6.2.2.3.7. Byte 112, Cabled PCIe Capabilities 2

Table 6-8 lists the bit descriptions for the PCI Express cable port capabilities register 2

**Table 6-8. PCI Express Cable Port Capabilities Register 2**

| Bit | Value | Description |
|---|---|---|
| 0-2 | 000b | X1 |
| | 001b | X2 |
| | 010b | X4 |
| | 011b | X8 |
| | 100b | X12 |
| | 101b | X16 |
| | | All others Reserved |
| 3-7 | | RsvdP |

#### 6.2.2.3.8. Bytes 119-126 Password Entry and Change

Bytes 119-126 are reserved for an optional password entry function. Password entry bytes are write only and are retained until power down, reset, or rewritten by the Fixed-side.

Cable vendors are permitted to use this function to implement write protection of Serial ID and other read only information. Password access must not be required to access Free-side device data in either the lower memory page 00h or upper page 00h. Note that multiple manufacturer passwords are permitted to be defined to allow selective access to read or write to various sections of memory, as allowed above.

Fixed-side Subsystem manufacturer and cable manufacturer passwords must be distinguished by the high order bit (bit 7, byte 123). All Fixed-side Subsystem manufacturer passwords must fall in the range of 0000 0000h to 7FFF FFFFh, and all cable manufacturer passwords in the range of 8000 0000h to FFFF FFFFh. Fixed-side Subsystem manufacturer passwords must be initially set to 0000 1011h in new cables.

Fixed side system manufacturer passwords are permitted to be changed by writing a new password in Bytes 119-122 when the correct current fixed side manufacture password has been entered in 123-126; with the high order bit being ignored and forced to a value of 0 in the new password. The password entry field must be set to 00000000h on power up and reset.

Page 02 is optionally provided as user writable EEPROM. The fixed side reads or writes this memory for any purpose. If Page 00h Byte 129 bit 4 is set, however, the first 10 bytes of Page 02h Bytes 128-137 are used to store the CLEI code for the free side device.

Common Language Equipment Identification is SFF-8636.

#### 6.2.2.3.9. Byte 127 Page Select

Byte 127 is used to select the upper page. A value of 00h indicates upper memory page 00h is available to be mapped to locations 128 to 255. All other values are reserved for future use.

#### 6.2.2.3.10. Byte 147, Cable Technology

The Cable Technology register indicates to the Fixed-Side Subsystem what type of cable is inserted. This information is used in conjunction with Bytes 186-189 to determine optimal equalizer settings for transmitters and receivers. FSS designers should see section 4.2.1.1 to understand how to interpret that data in Table 6-9 and use it for configuration of the PCI-Express devices.

**Table 6-9. Cable Technology**

| Value | Description |
|---|---|
| 0h | Active Optical Cable |
| 8h | Undefined/Other |
| Ah | Copper Cable, Unequalized |
| Bh | Copper Cable, Passively Equalized |
| Ch | Copper Cable, TX and RX Limiting Active Equalizers |
| Dh | Copper Cable, Rx Limiting Active Equalizer |
| Eh | Reserved |
| Fh | Copper Cable, Active Linear Equalizer |
| | All others Reserved |

#### 6.2.2.3.11. Bytes 148-163, Vendor Name

The Vendor Name is a 16-byte field that contains ASCII characters, left-aligned, and padded on the right with ASCII spaces (20h). The Vendor Name must be the full name of the corporation, a commonly accepted abbreviation of the name of the corporation, or the stock exchange code for the corporation. A value of all zeroes in the field indicates that the Vendor Name is unspecified. At least one of the Vendor Name or the PCI-SIG Vendor ID must contain valid non-zero data.

#### 6.2.2.3.12. Bytes 165-166, PCI-SIG Vendor ID

This is the 2-byte Vendor ID, allocated by the PCI-SIG. A value of all zeroes indicates that the Vendor ID is unspecified.

#### 6.2.2.3.13. Bytes 168-183, Vendor Part Number

The Vendor Part Number is a 16-byte field that contains ASCII characters, left aligned and padded on the right with ASCII spaces (20h). This field defines the vendor part number or product name. A value of all zeroes in the field indicates that the Vendor Part Number is unspecified.

#### 6.2.2.3.14. Bytes 184-185, Vendor Revision

The Vendor Revision is a 2-byte field that contains ASCII characters, left aligned, and padded on the right with ASCII spaces (20h) and defines the vendor's product revision number. A value of all zeroes in the field indicates that the Vendor Revision is unspecified.

#### 6.2.2.3.15. Bytes 186-189, Attenuation at Frequency

The data in these fields indicate the loss of the cable assembly and mated connectors in increments of 1 dB for each frequency. Test fixture loss must not be included in this data. A value of all zeroes indicates that the attenuation is not known or is unavailable. It is recommended that values of all zeros be used for cables in which the cable isolates the electrical waveform from either of the fixed sides. Examples of this include active optical cables (Cable Technology field value of 2xh) and cable with limiting amplifiers at both the Tx and Rx ends of the cable (Cable Technology field value of 6xh).

#### 6.2.2.3.16. Byte 190, Maximum Case Temperature

Max case temperature in degrees Celsius. The value in degrees is translated directly to a hex value. For an unspecified value, 00h, a limit of 70d C is assumed. A value representing zero degrees is not allowed.

**Note:** An unspecified value is considered 00h and this is interpreted as 70d C rather than 0d C.

#### 6.2.2.3.17. Byte 191, Checksum-Base

A check code to verify the first 63 bytes of the upper memory page. The value is the low order eight bits of the sum of the contents of bytes 128-190.

#### 6.2.2.3.18. Byte 195, Options

This field is reserved for future use to indicate various cable options including the use of Page 1 and 2 for the upper page.

#### 6.2.2.3.19. Bytes 196-211, Vendor Serial Number

The Vendor Serial Number is a 16-byte field that contains ASCII characters, left aligned, and padded on the right with ASCII spaces (20h) and defines the vendor's serial number for the Free-side device. A value of all zeroes in the field indicates that the Vendor Serial Number is unspecified.

#### 6.2.2.3.20. Bytes 212-217, Vendor Date Code

A six-byte field containing the vendor's date code in ASCII. Bytes 212-213 contain the two year digits (00=2000, 99=2099). Bytes 214-215 contain the month (01=Jan, 12=Dec). Bytes 216-217 contain the day of the month (01-31).

#### 6.2.2.3.21. Bytes 218-219, Vendor Lot Code

Vendor defined lot code. This is permitted to be all zeros.

#### 6.2.2.3.22. Byte 223, Checksum-Extended

A check code to verify the 31 bytes above the Checksum-Base in the upper memory page. The value is the low order 8 bits of the sum of the contents of bytes 192-222.

#### 6.2.2.3.23. Bytes 224-225, Vendor Specific Information

Vendor specified data.

## 6.2.3. Non-Volatile Memory Specification

The memory transaction timings are given in Table 6-10.

**Table 6-10. Non-volatile Memory Specifications**

| Parameter | Symbol | Min | Max | Unit | Conditions |
|---|---|---|---|---|---|
| Complete Single or Sequential Write | $t_{WP}$ | | | ms | Complete (up to) 4-byte Write |
| Serial Interface Clock Hold off (Clock Stretching) | $t_{SCL.HOLDOFF}$ | | See $T_{CMI\_CLK\_HOLD}$ in Table 2-2 | µs | Maximum time the cable module holds the CMICLK line low before continuing with a read or write operation |
| Endurance (Write Cycles) | | 50,000 | | Cycles | 70º C |

# 6.2.4. Cable Electrical Performance

**Note:** Values and limits for frequency domain parameters are typically listed in this document as Loss Values, and are therefore positive. However, since they are losses, the actual measured values, either simulated or measured, would be negative. Table 6-11 lists the electrical performance requirements for PCI Express external cables.

Cable assembly compliance testing is done using frequency domain techniques. The insertion loss, return loss, and crosstalk requirements are provided starting in Section 6.2.4.1. Cable assembly time domain testing is be performed against the eye diagrams in lieu of the frequency domain parameters. This allows trade-offs between the cable assembly interconnect parameters and implementations such as equalization. Table 6-11 lists a summary of the cable assembly differential characteristics and is included for convenience. The referenced section contains the specification requirement.

## Table 6-11. Cable Assembly Differential Characteristics Summary

| Description | Values for 8.0 GT/s[1] | Unit | Reference |
|---|---|---|---|
| Maximum Insertion Loss at 8.0 GHz<br>Maximum Insertion Loss at 4.0 GHz<br>Maximum Insertion Loss at 2.5 GHz<br>Maximum Insertion Loss at 1.25 GHz<br>Maximum Insertion Loss at 625 MHz | 14.6<br>9.6<br>7.3<br>4.9<br>3.3 | dB | See section 6.2.4.1 |
| Minimum Return Loss at 8.0 GHz<br>Minimum Return Loss at 4.0 GHz<br>Minimum Return Loss at 2.5 GHz<br>Minimum Return Loss at 1.25 GHz | TBD<br>TBD<br>TBD<br>TBD | dB | See section 6.2.4.2 |
| Minimum Return Loss at 8.0 GHz<br>Minimum Return Loss at 4.0 GHz<br>Minimum NEXT Loss at 2.5 GHz<br>Minimum NEXT Loss at 1.25 GHz | TBD<br>TBD<br>TBD<br>TBD | dB | See section 1.1.1.1 |
| Minimum Return Loss at 8.0 GHz<br>Minimum Return Loss at 4.0 GHz<br>Minimum MDNEXT Loss at 2.5 GHz<br>Minimum MDNEXT Loss at 1.25 GHz | TBD<br>TBD<br>TBD<br>TBD | | See section 6.2.4.3 |
| Minimum Return Loss at 8.0 GHz<br>Minimum Return Loss at 4.0 GHz<br>Minimum ELFEXT at 2.5 GHz<br>Minimum ELFEXT at 1.25 GHz TBD | TBD<br>TBD<br>TBD<br>TBD | | See section 6.2.4.4 |
| Minimum Return Loss at 8.0 GHz<br>Minimum Return Loss at 4.0 GHz<br>Minimum MDELFEXT at 2.5 GHz<br>Minimum MDELFEXT at 1.25 GHz | TBD<br>TBD<br>TBD<br>TBD | | See section 6.2.4.6 |
| Minimum DCMC below 2.5 GHz | TBD | | See section 6.2.4.9 |

**Notes:**
1. Applies to copper cable assemblies without any form of non-linear equalization.
2. Values must only include the cable assembly and mated connectors.

## 6.2.4.1. Maximum Insertion Loss

### 6.2.4.1.1. 8.0 GT/s

The insertion loss for each pair in a cable assembly operating at 8.0 GT/s must meet the values using Equation 6-1.

**Equation 6-1. 8.0 GT/s Maximum Insertion Loss**

$$InsertionLoss(f) \le 0.108\sqrt{f} + 0.0005 \times f \quad dB \quad 100 \le f \le 12000$$

where $f$ is the frequency in MHz.

Controlling the ISI jitter component through additional equalization within the cable assembly is allowed. Details of such implementations are beyond the scope of this specification. Figure 6-4 shows a graph of cable assembly insertion loss.

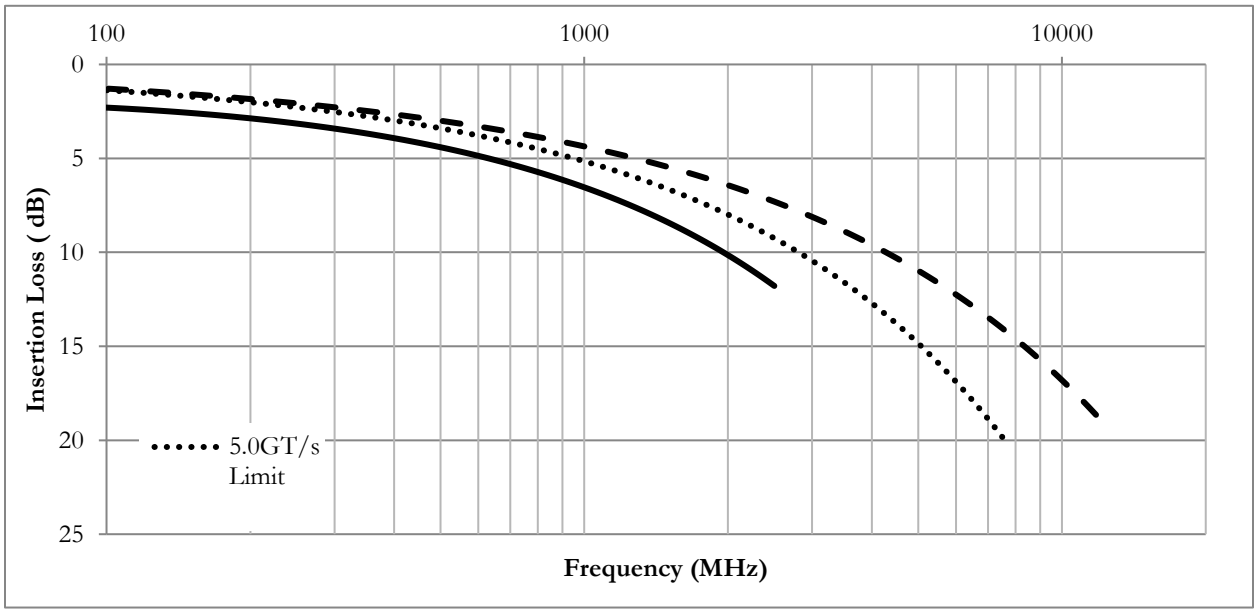


**Figure 6-4. Cable Assembly Insertion Loss (Needs updating in 0.9 Draft)**

## 6.2.4.2. Minimum Return Loss

### 6.2.4.2.1. 8.0 GT/s

The return loss of each pair of the cable assembly must meet the values determined using Equation 6-2.

**Equation 6-2. 8.0 GT/s Minimum Return Loss**

$$ReturnLoss(f) \ge 12 \quad dB \qquad 100 \le f(MHz) \le 2500$$

$$ReturnLoss(f) \ge 6.5 \quad dB \qquad 2500 < f(MHz) \le 8000$$

where $f$ is the frequency in MHz.

## 6.2.4.3. Minimum Multiple Disturber NEXT Loss

The NEXT loss that is coupled into a receive lane is from multiple transmit lanes at the near-end. To ensure the total NEXT loss coupled into a receive lane is limited; a multiple disturber NEXT loss is specified.

### 6.2.4.3.1. 8.0 GT/s

The MDNEXT loss between a receive lane and all transmit lanes (e.g., closest in proximity) must meet the values determined using Equation 6-3.

**Equation 6-3. 8.0 GT/s Minimum MDNEXT Loss**

$$MDNextLoss(f) \geq 33 - 14 \times \log\left(\frac{f}{4000}\right) \quad dB \quad 100 \leq f \leq 8000$$

where $f$ is the frequency in MHz.

MDNEXT loss is determined by summing the adjacent individual pair-to-pair differential NEXT loss values using Equation 6-4.

**Equation 6-4. 8.0 GT/s MDNEXT Loss Definition**

$$MDNextLoss(f) \text{ is defined as } -10 \times \log\left(\sum_{i=1}^{n} 10^{-NextLoss(f)_i/10}\right) \quad dB$$

where $f$ is the frequency in MHz.

$NextLoss(f)_i$ is the magnitude in dB of the NEXT loss for frequency $f$ for each measurement *i*, and *n* is the number of measurements. If the pair-to-pair NEXT loss measured between 100 MHz and 2500 MHz is greater than 50 dB over the entire frequency range, that measurement sweep does not need to be included in the MDNEXT summation.

## 6.2.4.4. Minimum ELFEXT

Equal level far-end crosstalk (ELFEXT) is specified to limit the far-end crosstalk (FEXT) appearing at a Receiver at the far end of a lane (disturbed lane) which is coupled from another lane (disturbing lane) with the noise source (Transmitters) at the near end.

ELFEXT is determined from the differential pair-to-pair FEXT and the insertion loss of the disturbed pair using Equation 6-5.

**Equation 6-5. ELFEXT Determined from the Differential Pair-to-Pair FEXT**

$$ELFEXT(f) \text{ is defined as } FEXT(f) - InsertionLoss(f) \quad dB$$

where $f$ is the frequency in MHz.

## 6.2.4.5. Minimum Multiple Disturber ELFEXT

Equal level far-end crosstalk (ELFEXT) is specified to limit the far-end crosstalk (FEXT) appearing at a Receiver at the far end of a lane (disturbed lane) which is coupled from another lane (disturbing lane) with the noise source (Transmitters) at the near end.

The FEXT loss that is coupled into a receive lane is permitted to be from multiple transmit lanes. To ensure the total FEXT loss coupled into a receive lane is limited; a multiple disturber ELFEXT loss is specified.

The MDELFEXT between a receive lane and all transmit lanes (e.g., closest in proximity) must meet the values determined using Equation 6-6.

### Equation 6-6. Determining MDELFEXT by Summing Individual Pair-to-Pair ELFEXT

$$MDELFEXT(f) \text{ is defined as } -10 \times \log\left(\sum_{i=1}^{n} 10^{-ELFEXT(f)_i/10}\right) \quad dB$$

where $f$ is the frequency in MHz.

$ELFEXT(f)_i$ is the magnitude in dB of the ELFEXT at frequency $f$ for each measurement *i*, and *n* is the number of measurements

If the pair-to-pair ELFEXT measured between 100 MHz and 2500 MHz is greater than 40 dB over the entire frequency range, that measurement sweep does not need to be included in the MDELFEXT summation

### 6.2.4.5.1. 8.0 GT/s

The MDELFEXT between a receive lane and all transmit lanes (i.e., closest in proximity) must meet the values determined using Equation 6-7.

### Equation 6-7. 8.0 GT/s MDELFEXT Between a Receive Lane and all Transmit Lanes

$$MDELFEXT \quad Loss(f) \geq 30 - 22.8 \times \log\left(\frac{f}{4000}\right) \quad dB \quad 100 \leq f \leq 8000$$

where $f$ is the frequency in MHz.

### 6.2.4.6. Maximum Pair-to-Pair Skew

The difference in propagation delay, or skew, between all cable assembly pair combinations, must not exceed the cable skew limit $S_C$ in Section 3.2.4.

### 6.2.4.7. Maximum Intra-Pair Skew

#### 6.2.4.7.1. 8.0 GT/s

There is not an explicit intra-pair skew requirement for 8.0 GT/s.

### 6.2.4.8. Maximum DC Resistance Single-Ended Sideband

The resistance of any conductor used for single-ended Sideband signaling must not exceed 10 Ω.

### 6.2.4.9. Differential-to-Common Conversion Loss

The minimum differential-to-common mode signal conversion loss in the cable assembly must be determined using Equation 6-8. This specification parameter only applies to cable assemblies operating at 8.0 GT/s.

**Equation 6-8. 8.0 GT/s Differential-to-Common Conversion Loss**

$$DCMC(f) \geq 18dB \quad 100 \leq f \leq 12000$$

where $f$ is the frequency in MHz

### 6.2.4.10. CMI Line Capacitance

The total capacitance for a CMI line within the completed cable assembly is defined by $C_{CMI\_CABLE\_LOAD}$. If the cable assembly bus capacitance is over $C_{CMI\_CABLE\_LOAD}$, active devices are employed, but all the AC specifications must still be met. For impedance targets of the CMI lines, see Table 2-2. The logic threshold requirements that must be met by active devices in the cable are in Table 2-1.

### 6.2.4.11. HiPot Requirement

The minimum HiPot requirement for cable assemblies must be 240 VDC for 100 ms.

## 6.3. Shielding

All cable assemblies must provide 360 degree shielding from end-to-end. EMI gaskets are required at the mating interface to reduce EMI emissions to an acceptable level.

7

# 7. Other Considerations

When expanding a system using a cabled interface, there are platform considerations that system vendors and users need to be aware of. While individual links function, the system must deal with configurations that have been beyond the scope of perceived usage models. Items such as bus enumeration, memory allocations, and platform timing need to be understood for the entire system to work correctly. Some of these considerations are presented here for awareness.

## 7.1. Discovery of PCI Express Topology

A cabled interface to an existing system allows for a sizable increase in the number of system components. In the case of a cabled PCI Express interface, the number of busses grow to approach the limit of the specification. Some system firmware are not designed to handle this seamlessly and is a problem for the host system. System integrators and end users should be aware that it is possible to exceed what the system's original firmware was designed to operate with, which leads to unknown system behaviors.

## 7.2. Memory Allocation

Adding additional devices to an existing system allows for a sizable increase in the number of system components, and therefore an increase in the system resources required to enable these components to operate. A key component of this is system memory. PCI devices are allocated memory upon discover, and requests more. A cabled system allows for an extensive number of devices to be added to the original system and the system hardware and firmware is not be able to predictably operate with demands larger than what the original system was designed for.

## 7.3. Completion Timeouts of PCI-Express Transactions

Adding additional devices to an existing system allows for a sizable increase in the number of system components. These components are added in a cascaded fashion through a series of PCI Express switches. There are aspects to the protocol that require operations to complete in an expected timeframe. While the standard does allow for large delays, and programmable delays in many implementations, these flexibilities are not required by the standard. This is a problem for the host system. System integrators and end users should be aware that exceeding what the system's original firmware was designed to operate with leads to unknown system behaviors.

## 7.4. Bus Performance and Throughput for Interfaces with Long Flight Times

An active cable, either copper or fiber, allows for a significant propagation delay of a signal between devices; well beyond what the base specification originally intended. While the protocol deals with long delays, there are impacts to system performance. The protocol includes flow control credits to limit the amount of data in flight as to not overflow a device's link receiver buffers. Some data transactions must send an acknowledgement that data was received before more data transactions are sent. Independent of device performance, the delay induced by the physical transmission medium is at least twice the propagation delay (once for the data transmission and again for the return acknowledgement) before the next data transaction is sent. The protocol has some built-in checks to make sure a data transaction is not unwittingly lost after it is sent. In addition to the Update FC being returned for flow control, the ACK must be returned on a finite schedule. If the data is not received within a certain window, the data transaction is sent again, or replayed. The protocol allows for data transaction to be replayed three times, for a total of four attempts at data transmission and acknowledgement. After the fourth attempt, a correctable error occurs, and the link attempts to recover itself.

## 7.5. Interoperability with SAS cables

The ability to use off-the-shelf cable assemblies that are compliant with SFF-8449 and SFF-8636. However, the link setup information contained in the cable assembly memory device are slightly different than that of a PCIe specific cable assembly.

An SFF-8449 cable assembly has a $T_{MGMT\_INIT}$ time of 2000 ms.

**Note:** SFF-8449 pin definitions are different than PCI Express definitions, though the cables are functionally compatible. Refer section 5.2.1 for further details.

The CMI controller in the Fixed-Subsystems must translate the SFF-8636 data into a format that is useful for programming the PCI-Express device. The controller determines if the cable is PCI-Express by reading byte 111, bit 0.

It should be noted that SFF-8449 cables do not transmit sideband messages and any behaviors that depend on sideband messages when this type of cable is used to create the cabled link needs to be adjusted. Examples include boot sequencing and local resets.

Key differences in the memory map relating to configuration data are:

- Bytes 108-109
  In SFF-8636, initially developed for SAS, Bytes 108-109 contain propagation delay in multiples of 10 ns. However, if Byte 111 is set (equal to 01b), indicating a PCI Express application, Bytes 108-109 contain propagation delay and lane-to-lane cable skew in multiples of 250 ps, as defined in Section 6.2.2.3.5.
- Bytes 111-112
  Bytes are all zeroes, per section 6.1, SFF-8644 Cable Assembly of the common memory map specification.
- Byte 147
  PCI-Express uses a subset of the options defined in an SFF-8449 cable. Values of 8h, Ah, Bh, Ch, Dh, and Fh are the same. In an SFF-8449 cable, a value of 0h specifies a particular type of optical driver. PCI-Express is not required to distinguish between types of optical drivers and the value of 0h indicates an optical driver is used. Values of 1h-7h and 9h all indicate other types of optical drivers. Any settings beyond what is recommended in chapter 4 is beyond the scope of this specification.
- Bytes 165-166 are not to be PCISIG ID
- Byte 186-189
  Indicate performance characteristics of the cable. Detailed recommendations to implement PCI-Express devices settings are beyond the scope of this specification. However, differences are highlighted to allow firmware designers to implement solutions if they choose to.
  - Byte 186
    Specify either wavelength or copper attenuation. Copper attenuation specified at 0 GHz if Byte 147 is Ah-Fh and Byte 133, bit 6, set to 1b. Byte 147 indicates if this is a copper or optical interface. If it is optical, this field is not relevant for configuration. If it is copper, this is the bit rate, not the Nyquist frequency, and is likely to indicate more loss than what is present at Nyquist.
  - Byte 187 specify either wavelength or copper attenuation. Copper attenuation specified at 0.0 GHz if Byte 147 is Ah-Fh and Byte 133, bit 6, set to 1b. Byte 147 indicates if this is a copper or optical interface. If it is optical, this field is not relevant for configuration. If it is copper, this is the bit rate, not the Nyquist frequency, and is likely to indicate more loss than what is present at Nyquist.
  - Byte 188 specify either wavelength or copper attenuation. Copper attenuation specified at 7 GHz if Byte 147 is Ah-Fh and Byte 133, bit 6, set to 1b. Byte 147 indicates if this is a copper or optical interface. If it is optical, this field is not relevant for configuration. If it is copper, this frequency is a 12% below the bitrate, or 75% above the Nyquist frequency.

Byte 189 specify either wavelength or copper attenuation. Copper attenuation specified at 12 GHz if Byte 147 is Ah-Fh and Byte 133, bit 6, set to 1b. If it is copper, this is the bit rate, not the Nyquist frequency, and is likely to indicate more loss than what is present at Nyquist. If it is copper, this frequency is a 25% below the bitrate, or 50% above the Nyquist frequency.

1580

A

# 1581 Appendix A. Acknowledgements

1582 The following persons were instrumental in the development of the *PCI Express External Cabling*
1583 *Specification* (Company affiliation listed is at the time of specification contributions):

1584